A Canon Communications LLC Publication

EDN®

VOICE OF THE ENGINEER

AUG 12
Issue 15/2010
www.edn.com

IP INTEGRATION: IS IT THE REAL SYSTEM-LEVEL DESIGN?

Page 34

MEASURING WIDEBAND-AMPLIFIER SETTLING TIME

Page 24

DATA STORAGE YIELDS INCREASED DESIGN PRODUCTIVITY

Page 41

Thermal Challenge: VRM

High Ambient / High Watt Dissipation / Low Air Flow

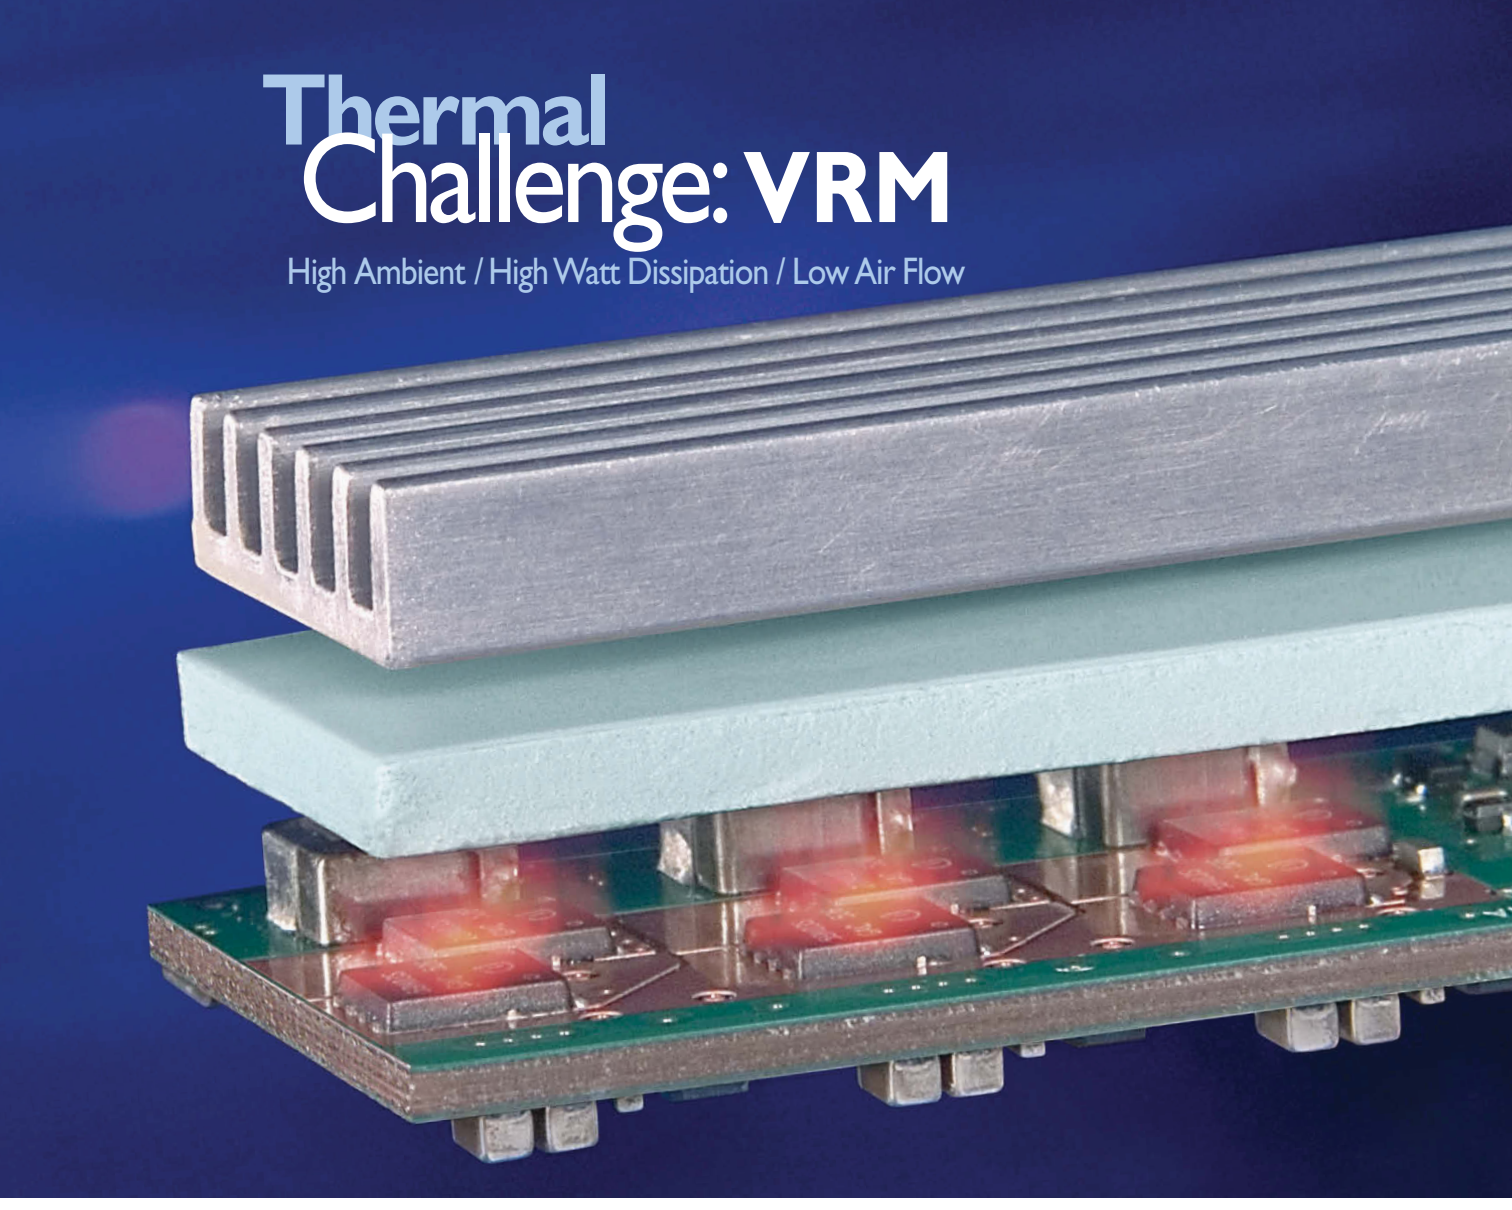

Solution: New Bergquist S-Class Gap Pad® 5000S35

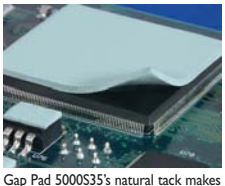

Gap Pad 5000S35's natural tack makes application clean and easy to handle.

Gap Pad S-Class is perfectly suited for high performance applications such as VRMs, BGAs and ASICs.

Bergquist's newest S-Class is the perfect combination of softness, low thermal resistance and high thermal conductivity. With a low bulk hardness (35 Shore 00) and high thermal conductivity (5.0 W/m-K) it conforms to demanding contours while maintaining its structural integrity. It is an ideal gap filling solution for applications with fragile components that can be damaged by harder materials that cause higher mounting pressure on components. Gap Pad 5000S35 is also an excellent solution for DVD drives, memory modules, and PC boards to chassis.

Excellent interfacing and wet-out makes GP 5000S35 ideal for fragile components with demanding contours and stack-up tolerances.

Ultra soft S-Class integrity offers easy application.

Gap Pad 5000S35 has a natural tack that eliminates the need for additional adhesive layers that can inhibit thermal performance. Its super soft, yet elastic nature provides excellent interfacing and wet-out, even to surfaces with high roughness or topography. Gap Pad 5000S35 features an embedded-fiberglass reinforcement that makes it puncture, shear and tear resistant. No tearing, flaking or crumbling – just clean and easy handling during the assembly process.

FREE S-Class swatch kit.

Visit our web site or call to qualify for your FREE S-Class swatch kit with product details.

FREE Gap Pad S-Class Swatch Kit

Call **1.800.347.4572** or visit **www.bergquistcompany.com/coolrun**

THE BERGQUIST COMPANY

18930 West 78th Street • Chanhassen, MN 55317 • A ISO9001:2000 registered facility
(800) 347-4572 • Phone (952) 835-2322 • Fax (952) 835-0430 • www.bergquistcompany.com

Thermal Materials • **Thermal Substrates** • **Fans and Blowers**

TAKE CONTROL

From Prototype to Production!

SAME-DAY SHIPPING on orders entered by 8:00 pm CENTRAL
NEXT-DAY DELIVERY

1. "Always has available product"
2. "Has an excellent website"
3. "Supports your global distribution needs"
4. "Provides on time delivery"
5. "Is easy to do business with"
6. "Most preferred distributor"

Digi-Key rated #1 in all categories.
Source: *TechInsights*, Distributor Evaluation Study, May 2009.

Digi-Key® CORPORATION

Over 465,000 electronic components in stock from more than 440 supplier partners.*

The industry's broadest product selection available for immediate delivery

www.digikey.com

1.800.344.4539

*Digi-Key is an authorized distributor for all supplier partners. New products added daily. © 2010 Digi-Key Corporation, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA

NEDA Supplier Authorized Distributors

Windows® Life without Walls.™ Dell recommends Windows 7.

Let your vision take shape.

Image provided by Autodesk and created with Autodesk Revit software for building information modeling (BIM).

Unleash your creativity and imagine the possibilities. Then watch as they become reality with the power of Dell Precision™ workstations and Autodesk® BIM software.

Dell Precision™ workstations deliver the performance and graphics needed to run demanding applications with ease. Now your team can use data-rich modeling to evaluate new design options, predict building performance and communicate more productively.

- Scalable options – select systems are available with up to 192GB of memory* and 7.5TB of internal storage
- A full range of desktop, rack and mobile workstations to fit your needs
- ISV-certified for 95 leading applications, including Autodesk®
- Stay up and running with a 3-year Limited Hardware Warranty** and optional 24/7 Dell ProSupport™

Get Equipped ▸ **Find your ideal configuration online. Go to dell.com/smb/vision or call your sales rep at 1-877-952-3355**

Windows 7 Professional

Windows® 7 Professional operating system makes the things you do every day faster and easier. With fewer clicks, faster searching, easier browsing and simpler ways to connect, there's less between you and what you want to do.

*Actual capacity varies with preloaded material and operating environment and may be less. Significant system memory may be used to support graphics, depending on system memory size and other factors. ** For copy of Ltd Hardware Warranty, write Dell USA LP. Attn Warranties, One Dell Way, Round Rock, TX 78682 or visit dell.com/warranty. Remote Diagnosis is determination by online/phone technician of cause of issue, may involve customer access to inside of unit and multiple or extended sessions. If issue is covered by Limited Hardware Warranty and not resolved remotely, technician and/or part will be dispatched usually within 1 business day following completion of Remote Diagnosis. Availability varies. Other conditions apply.

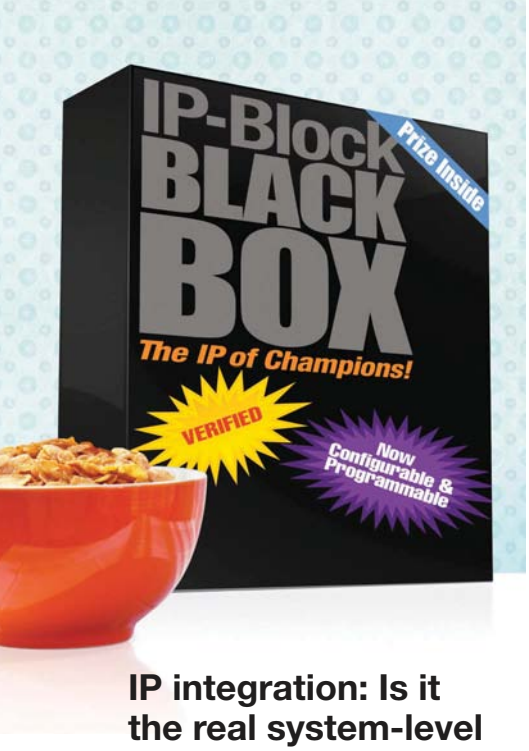

EDN 8.12.10 contents

IP integration: Is it the real system-level design?

By Ron Wilson, Executive Editor

Measuring wideband-amplifier settling time

By Jim Williams, Linear Technology

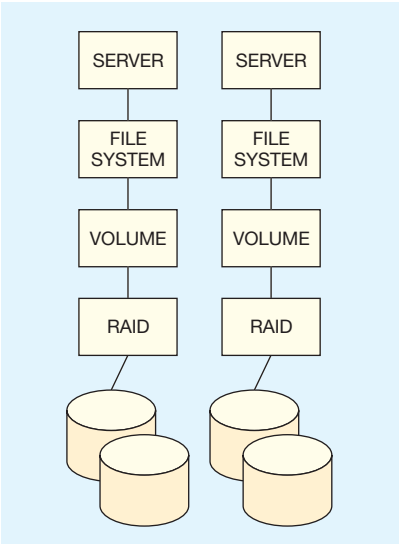

Data storage yields increased design productivity

By Paul Rutherford, Isilon Systems

pulse

DESIGNIDEAS

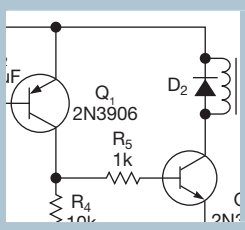

Yes, chip inductors are among Coilcraft's biggest sellers. But did you know we also offer a full range of other RF products like these?

They're all in stock for immediate shipment and all available as free samples for testing and evaluation.

For the full picture, visit **www.coilcraft.com/RF**

Broadband Conical Inductors

High impedance from 10 MHz to 40 GHz. Perfect for RF to millimeter-wave choking and bias tees. Available with leads or in surface mount versions for ruggedness and easy mounting.

Air Core Inductors

For the highest possible Q and current handling, nothing beats our air core "Spring" inductors. New models offer more inductance values and smaller size.

RFID Transponder Coils

A variety of antenna coils for 125 kHz systems. Automotive grade models for tire pressure and keyless entry applications.

Wideband Transformers

Our low insertion loss transformers come in a wide range of turns ratios. Use for baluns, isolation or impedance matching.

Some people only think of us for chip inductors.

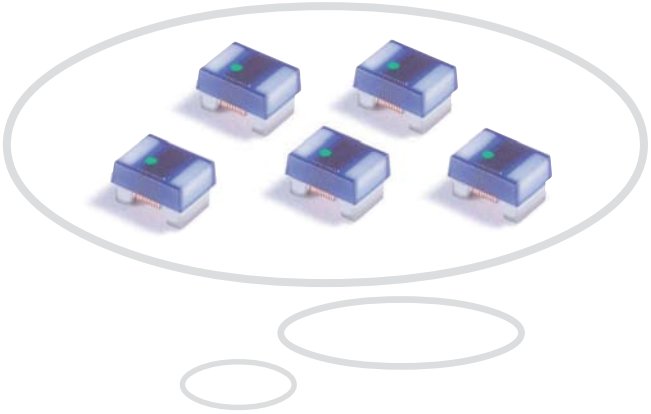

Coilcraft®

contents 8.12.10

54

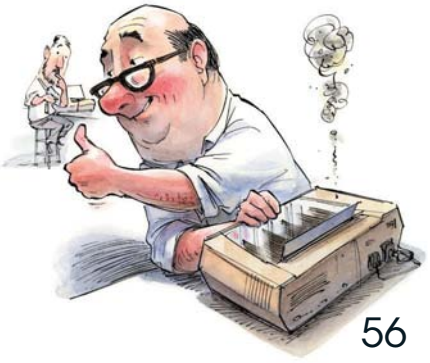
56

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out this Web-exclusive article:

Interoperable DRC/LVS language standard accelerates physical-verification turnaround time for advanced nodes
→www.edn.com/100812toca

FEED YOUR NEED

EDN's technical editors highlight notable new products, including analog and digital ICs, power components, sensors, passives, and boards. Contact the appropriate editor to submit products for consideration.
→www.edn.com/100812tocb

Readers' Choice

Here is a selection of recent articles and blog posts receiving high traffic on www.edn.com:

DC-voltage doubler reaches 96% efficiency
→www.edn.com/100812tocc

iTunes Store woes: Apple's arrogance steadily grows
→www.edn.com/100812tocd

Last hope for retail set-top boxes?
→www.edn.com/100812toce

The time is now for 3-D stacked die
→www.edn.com/100812tocf

Several key commodity components in "critically short supply"
→www.edn.com/100812tocg

EDN® (ISSN#0012-7515) is published semimonthly, 24 times per year, by Canon Communications LLC, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 310/445-4200; FAX 310/445-4299. Periodicals postage paid at Los Angeles, CA, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 763/746-2792. Email: EDN@kmpsgroup.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publisher cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. Canada Post: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to Pitney Bowes Inc, PO Box 25542, London, ON N6C 6B2. *Volume 55, Number 15 (Printed in USA).*

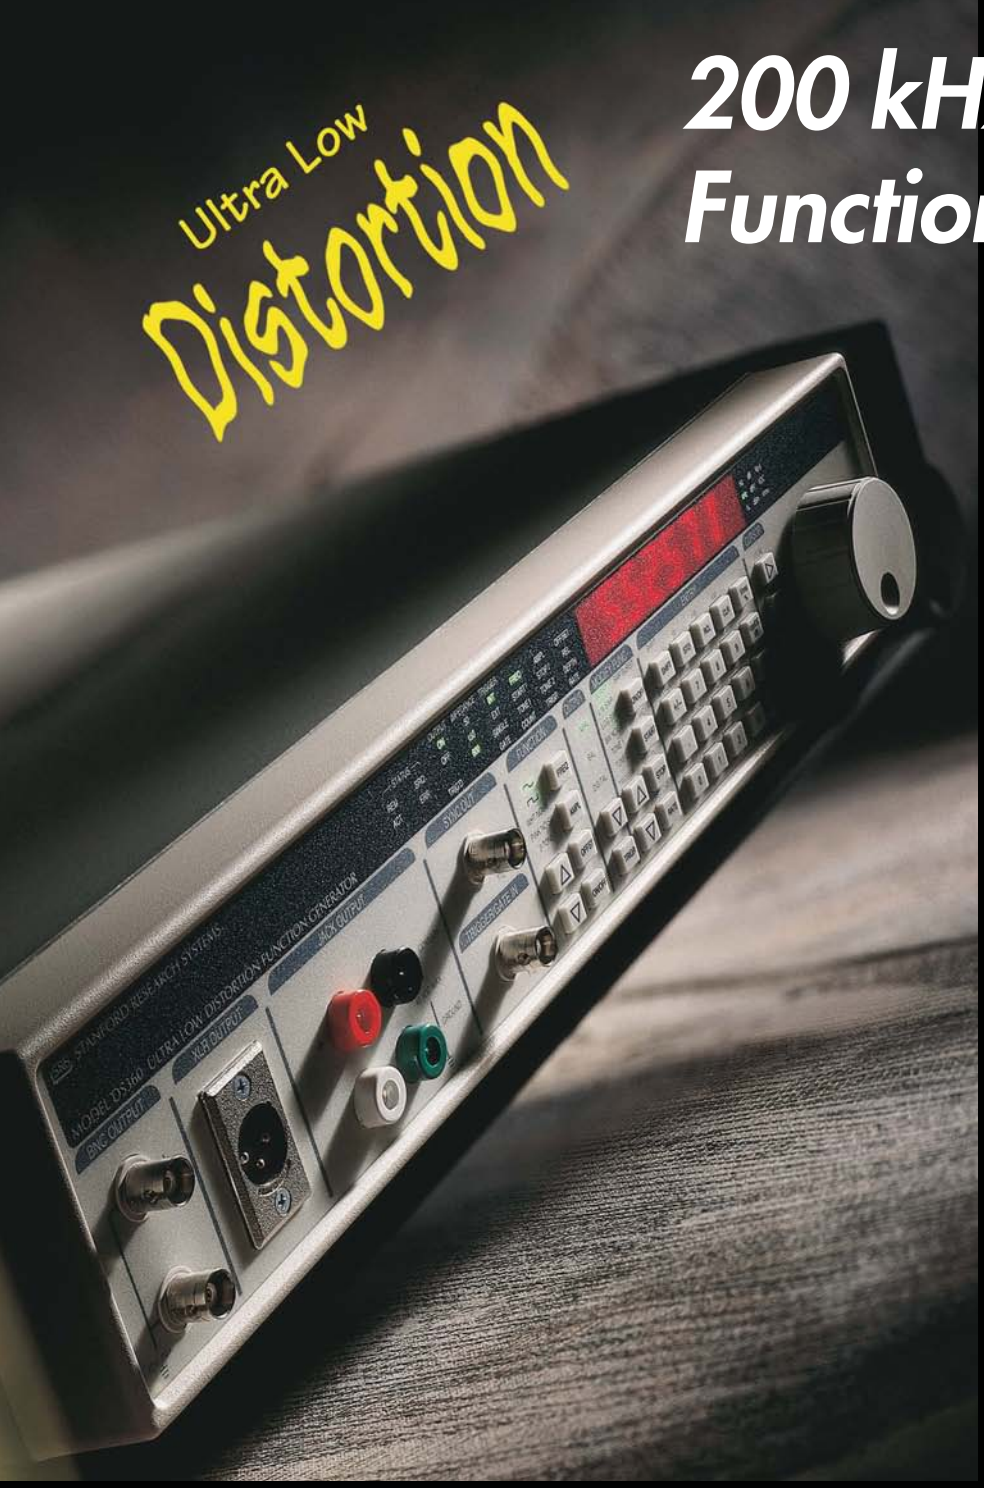

200 kHz DDS Function Generator

0.001 % THD
(–100 dB)

DS360...$2795 (U.S. list)

- 25 ppm freq. accuracy
- 1 mHz to 200 kHz
- 4-digit amplitude resolution
- Sine, square, two-tone, burst
- Pink and white noise
- Log & linear freq. sweeps
- 35 mVpp to 40 Vpp
- SPDIF/EIAJ and AES-EBU
- GPIB and RS-232 (std.)

The performance of a low-distortion signal source and the features of a digitally synthesized function generator are combined in a unique instrument, the DS360. With 0.001 % THD, 25 ppm frequency accuracy and functions that include sine waves, square waves, sweeps, and bursts, the DS360 delivers unparalleled precision and versatility. To make it complete we added two-tone capability (including SMPTE, DIM, CCIF), bandwidth limited white noise, SPDIF/EIAJ and AES-EBU digital outputs, and computer interfaces.

SRS **Stanford Research Systems**
Phone (408) 744-9040 www.thinkSRS.com

SUZANNE DEFFREE, MANAGING EDITOR, NEWS

India's $35 tablet: If it's possible, it could be a game changer

India is touting a $35 "cloud"-computing-tablet prototype that includes word-processing, Web-browsing, and videoconferencing functions. If $35 is too much for you, however, sit tight. Government officials are looking for partners to reduce the cost to approximately $10—less than I spend on coffee in a week and about the same price you would pay for a no-frills, pay-as-you-go cell phone at Target.

The move is part of a broader effort by India's government toward better education in the country and accompanies stated plans to bring broadband connectivity to the country's 25,000 colleges and 504 universities. Much like the OLPC (One Laptop Per Child) effort, which the Massachusetts Institute of Technology backs, India with its $35 tablet aims to provide a basic, low-cost mobile-computing device that will allow the country's citizens to engage in their own education and access the Internet, ultimately allowing for a higher quality of life.

Notably, India's tablet is much less expensive than the current $199 XO laptop from OLPC. Marvell in May announced that it had teamed with OLPC for a sleek $100 tablet, still at a price almost three times higher than that of India's $35 device.

India's tablet will most likely be Linux-based, although there has been no official word on software. The prototype mockup, exhibited in late July, shows a touchscreen and storage relying on memory cards rather than hard drives. It also has an extra-cost, according to some reports, solar-cell option, which could be the main source of power for some in rural India.

Information has not yet emerged on what components the tablet includes, what type of wireless it runs, or which manufacturer supplies the display. Details are also lacking about processor speed, battery life, compatibility, and display resolution. Promoting the device, India's MHRD (Ministry of Human Resource Development) also omitted details about how R&D will decrease the price and what companies will manufacture the tablet. Details notwithstanding, it's safe to say the tablet's performance and debut won't be as smooth as that of the iPad or any of the competing tablets from major OEMs for the established US, European Union, and Asia-Pacific markets, but they don't need to be. Anything is better than nothing in a country where many have no computing at all. And a cloud-based system allows á la carte personalization, a good match for low-income budgets.

You have to wonder which chip makers could afford to supply the tablet at such a low price point. Even a full $35 BOM (bill-of-materials) cost doesn't leave much wiggle room at current component prices, at least not for all the things India's MHRD is claiming the tablet will do. But would companies such as Intel or AMD want no part of such a tablet, which could be a game changer for the Indian market? Again, the tablet aims to bring the Internet to India, taking a first step toward computing for many in the low-income country. Chip makers may find themselves offering significant discounts to even nip at the massive market opportunity, especially when you look at the projected market size.

The US Census Bureau International Data Base estimates that India will be the most populous country by 2025, surpassing China, with its population continuing to expand and projected to exceed 1.65 billion people by 2050. Electronics executives now bend over backward to land a place in the Chinese market. Imagine how they will compete when India's access to Internet devices rises.

Some wonder whether the $35-tablet idea will fizzle out in the same way that India's previous promises for a $20 laptop did. Even if it never moves from prototype to consumer hands, however, there's something to consider: Developing nations don't need iPads; they don't even need our netbooks, and they would not be apt to pay the accompanying price tags. They may just need simple computing that accesses the Internet through a cloud-based system.

How will the electronics supply chain provide that access, assuming it chooses to serve the needs of developing markets rather than push products that meet the needs of developed nations? The answer to that question could force a breed of designs that see the electronics industry backtracking to simpler systems, possibly cutting into chip revenue. EDN

You can reach me at suzanne.deffree@cancom.com.

Our engineers needed a faster scope. A scope that would display 1 million waveforms per second. So we built one.

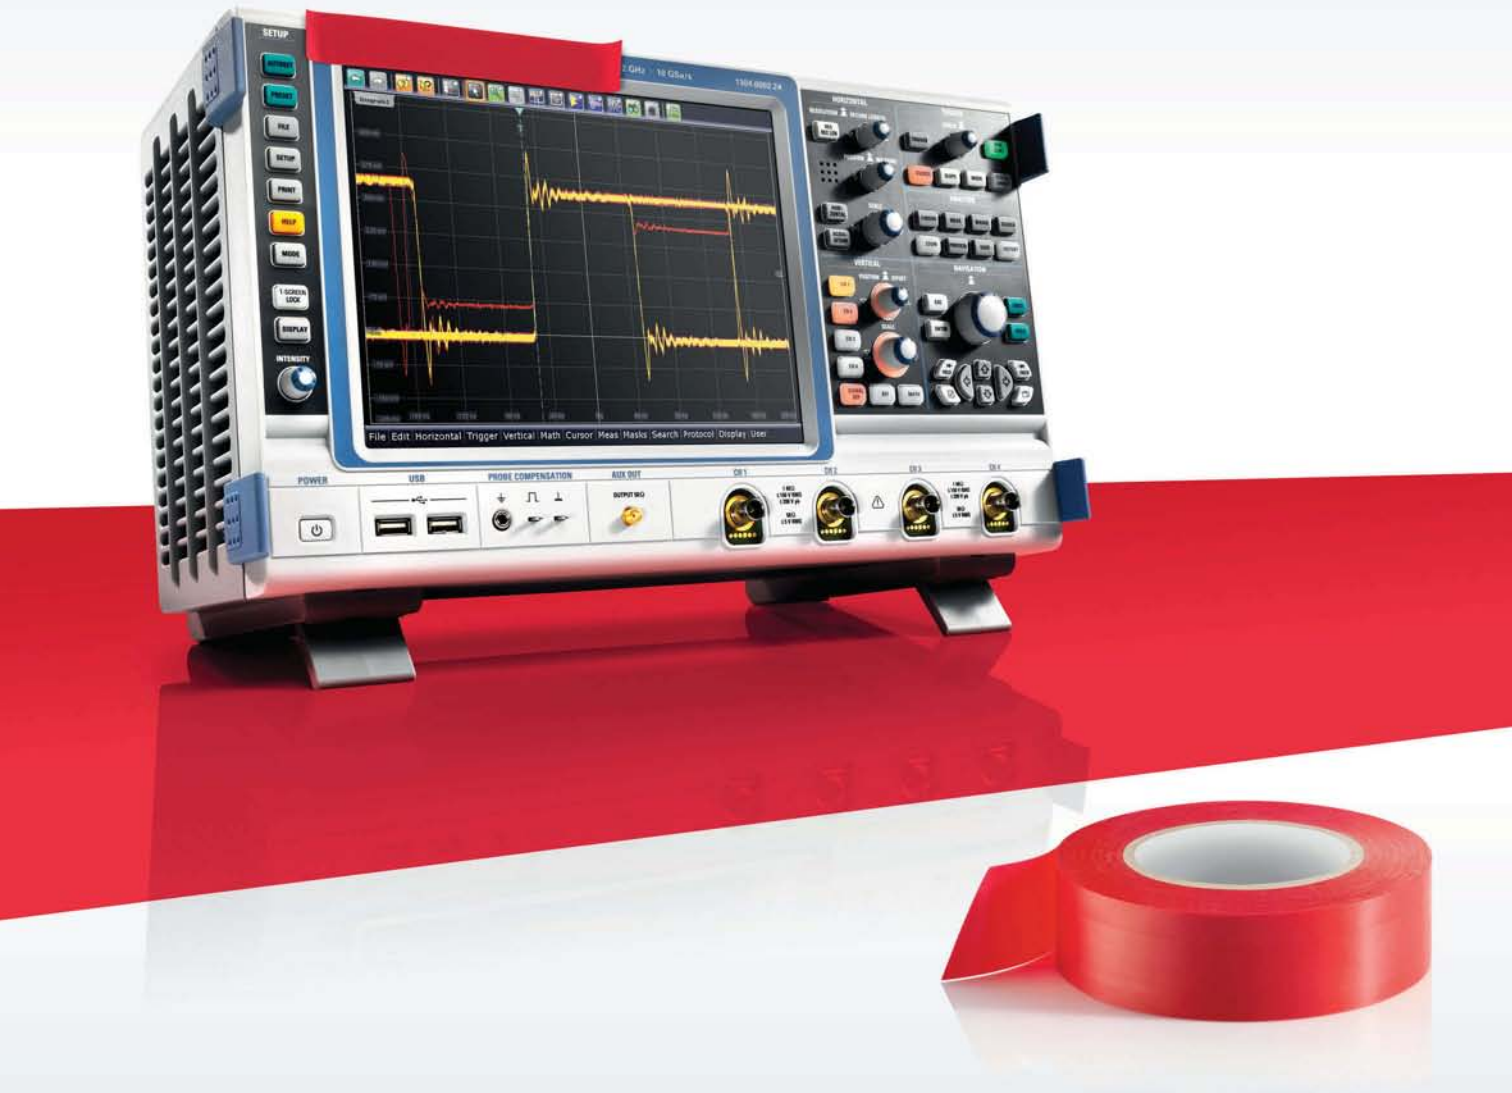

See for yourself

www.scope-of-the-art.com/ad/fasterscope

EDN

PUBLISHER, *EDN* WORLDWIDE
Russell E Pratt, 1-781-869-7982;
russell.pratt@cancom.com

ASSOCIATE PUBLISHER, *EDN* WORLDWIDE
Judy Hayes, 1-925-736-7617;
judy.hayes@cancom.com

EDITOR-IN-CHIEF, *EDN* WORLDWIDE
Rick Nelson, 1-781-869-7970;
richard.nelson@cancom.com

EXECUTIVE EDITOR
Ron Wilson

MANAGING EDITOR
Amy Norcross
1-781-869-7971;
fax: 1-781-862-4853;
amy.norcross@cancom.com
Contact for contributed technical articles

ANALOG
Paul Rako, Technical Editor
1-408-745-1994;
paul.rako@cancom.com

MASS STORAGE, MULTIMEDIA, PCs, AND PERIPHERALS
Brian Dipert, Senior Technical Editor
1-916-548-1225;
brian.dipert@cancom.com

NEWS
Suzanne Deffree, Managing Editor
1-631-266-3433;
suzanne.deffree@cancom.com

POWER SOURCES, ONLINE INITIATIVES
Margery Conner, Technical Editor
1-805-461-8242;
fax: 1-805-461-9640;
margery.conner@cancom.com

DESIGN IDEAS EDITOR
Martin Rowe,
Senior Technical Editor,
Test & Measurement World
edndesignideas@cancom.com

SENIOR ASSOCIATE EDITOR
Frances T Granville
1-781-869-7969;
fax: 1-781-862-4853;
frances.granville@cancom.com

ASSOCIATE EDITOR
Jessica MacNeil
1-781-869-7983;
jessica.macneil@cancom.com

CONSULTING EDITOR
Jim Williams, Staff Scientist,
Linear Technology
edn.editor@cancom.com

CONTRIBUTING TECHNICAL EDITORS
Dan Strassberg,
strassbergedn@att.net
Robert Cravotta,
robert.cravotta@embeddedinsights.com

COLUMNISTS
Howard Johnson, PhD, Signal Consulting
Bonnie Baker, Texas Instruments
Pallab Chatterjee, SiliconMap
Kevin C Craig, PhD, Marquette University

LEAD ART DIRECTOR
Marco Aguilera

ASSOCIATE ART DIRECTOR
Tim Burns

PRODUCTION
Michael Ciardiello,
Director of Premedia Technologies
Jeff Tade,
Production Director
Brian Wu,
Publications Production Manager
Jeff Polman, Derric Treece,
Senior Production Artists
William Baughman, Lucia Corona
Ricardo Esparza,
Production Artists

EDN EUROPE
Graham Prophet,
Editor, Reed Publishing
gprophet@reedbusiness.fr

EDN ASIA
Wai-Chun Chen,
Group Publisher, Asia
waichun.chen@cancom.com
Kirtimaya Varma,
Editor-in-Chief
kirti.varma@cancom.com

EDN CHINA
William Zhang,
Publisher and Editorial Director
william.zhang@cancom.com
Jeff Lu,
Executive Editor
jeff.lu@cancom.com

EDN JAPAN
Katsuya Watanabe,
Publisher
katsuya.watanabe@cancom.com
Ken Amemoto,
Editor-in-Chief
ken.amemoto@cancom.com

EXECUTIVE OFFICERS
Charles G. McCurdy,
Chief Executive Officer
Fred Gysi,
Chief Financial Officer
Mike Deering,
Chief Technology Officer
Ron Wall,
Senior Vice President, Publications
Kevin O'Keefe,
Senior Vice President, Events Division
Roger Burg,
Vice President, Operations
Jason Brown,
Vice President, E-media

EDN. 33 Hayden Avenue, Lexington, MA 02421. **www.edn.com**. **Subscription inquiries:** 1-763-746-2792; EDN@kmpsgroup.com. **Address changes:** Send notice promptly to PO Box 47461, Plymouth, MN 55447. Please provide an old mailing label as well as your new address. Allow two months for the change. Canon Communications LLC, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 1-310-445-4200; fax: 1-310-445-4299.

CANON COMMUNICATIONS LLC

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

Can light beams replace electronic signals in future computers?

Intel has announced an advance in the quest to use light beams to replace the use of electrons to carry data in and around computers. The company has developed a research prototype employing extremely thin and light optical fibers and representing the first silicon-based optical-data connection with integrated lasers. The link can move data over longer distances and do it many times faster than today's copper technology—as fast as 50 Gbps.

Justin Rattner, chief technology officer at Intel and director of Intel Labs, demonstrated the link at the Integrated Photonics Research Conference last month in Monterey, CA. The link, according to Rattner, is akin to a "concept vehicle" that allows Intel researchers to test new ideas and continue the company's quest to develop technologies that transmit data over optical fibers, using light beams from low-cost and easy-to-make silicon instead of costly and difficult-to-make devices using exotic materials, such as gallium arsenide. Although telecommunications and other applications use lasers to transmit information, current technologies are too expensive and bulky for PC applications.

"This achievement of the ... 50-Gbps silicon-photonics link with integrated hybrid silicon lasers marks a significant achievement in our long-term vision of 'siliconizing' photonics and bringing high-bandwidth, low-cost optical communications in and around future PCs, servers, and consumer devices," Rattner says.

The prototype is the result of a multiyear silicon-photonics research agenda. It comprises a silicon transmitter and a receiver chip, each integrating all the necessary building blocks from previous Intel technologies, including the first hybrid silicon laser that Intel co-developed with the University of California—Santa Barbara in 2006 as well as high-speed optical modulators and photodetectors that Intel announced in 2007.

The transmitter chip comprises four such lasers, whose light beams each travel into an optical modulator that encodes data onto them at 12.5 Gbps. The four beams then combine and output to a single optical fiber for a total data rate of 50 Gbps. At the other end of the link, the receiver chip separates the four optical beams and directs them into photodetectors, which convert data back into electrical signals.

Intel assembles both chips using low-cost manufacturing techniques. Company researchers are working to increase the data rate by scaling the modulator speed and increase the number of lasers per chip, providing a path to future terabit-per-second optical links.

—by Rick Nelson

▷**Intel Corp**, www.intel.com.

TALKBACK

"Just think of the fun that you would have if the system got out into the field and then a relatively normal voltage came along and popped the fuse, and it was one of those nice soldered-in fuses, and there was no spare." —Engineer William Ketel, in *EDN*'s Talkback section, at http://bit.ly/aDxWq0. Add your comments.

Intel's 50-Gbps silicon-photonics link allows researchers to develop and test new technologies that transmit data over optical fibers, using light beams from low-cost silicon instead of costly and hard-to-make devices using exotic materials, such as gallium arsenide.

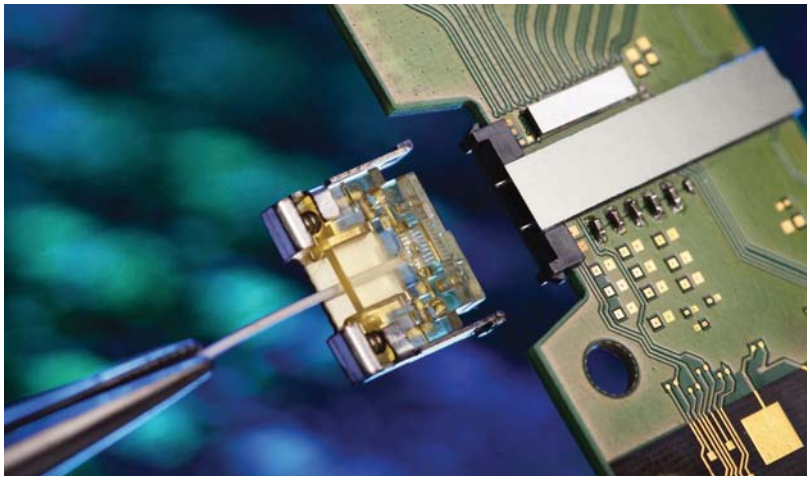

“I need a faster, more accurate audio analyzer.”

We knew you’d say that.

The new Agilent U8903A audio analyzer is four times more accurate and over seven times faster than its popular HP 8903B ancestor. In fact, it is among the fastest and most accurate currently in its class. It’s also a more fully capable audio analyzer, letting you make crosstalk measurements and phase measurements, FFT analysis, and graph sweep. That’s taking it to the next level. That’s Agilent.

Learn how to migrate effortlessly -
Free application note filled with tips
www.agilent.com/find/AgilentU8903A

u.s. 1-800-829-4444 canada: 1-877-894-4414

Agilent Technologies

30-MHz generators produce predefined functions, precision pulses, and low-jitter arbitrary waves

Agilent Technologies has added a pair of instruments to its 33000 Series of function/AWGs (arbitrary-waveform generators). The 33521A provides a single output channel, and the 33522A provides two channels. The generators produce sine, square, and pulse waveforms with repetition rates that you can set from 1 µHz to 30 MHz with 1-µHz resolution. The square waves have nominal rise and fall times of 8.4 nsec.

For less than $3000, the two-channel, half-rack-width 33522A delivers sine, square, and pulse waveforms to 30 MHz, plus a long list of built-in wave shapes and an unlimited number of arbitrary signals.

The pulse waveforms have rise and fall times that you can vary independently from 8.4 nsec to 1 µsec with 100-psec or three-digit resolution. The generators can produce pulse waveforms over the full sine- and square-wave repetition-rate range, whereas most competitors' units have a lower upper limit on pulse rate than they do on sine-wave frequency and square-wave repetition rate. Other built-in waveforms include positive and negative ramp, triangle, Gaussian noise with a repetition pattern that exceeds 50 years, PRBS (pseudorandom binary sequence), cardiac pulse, Gaussian pulse, exponential rise, exponential fall, haversine, Lorentz, time-dilated Lorentz, and sin(x)/x.

The new generators provide 1M-sample/channel memory depth for user-defined arbitrary waveforms, and a 16M-sample/channel memory depth is optional. The clock rate is 250M samples/sec, and the waveform definitions use 16 bits/sample, providing high time resolution, high vertical resolution, less than 0.04% total harmonic distortion, and less than 40-psec-rms jitter. Waveform sequencing allows point-by-point definition of waveforms that are even deeper than the physical memory. The instruments' proprietary architecture also overcomes a significant shortcoming of most function-synthesizer-based generators—high waveform-timing jitter under certain conditions—for example, when a waveform's duration is not an exact integer multiple of the sample clock period.

The generators fully comply with the LXI (local-area-network extensions for instrumentation) Class C Specification. They include USB (Universal Serial Bus) 2.0 and 10/100 BaseT Ethernet LAN ports for quick and easy connectivity to a PC or a network. A built-in Web page allows remote operation of the instruments from any Web browser. A GPIB (general-purpose-interface-bus) option allows you to connect the new generators with Agilent's 33210A, 33220A, and 33250A function/AWGs. Prices for the 33521A and 33522A start at $1930 and $2950, respectively.

—by Dan Strassberg

▷**Agilent Technologies**, www.agilent.com/find/33500.

DILBERT By Scott Adams

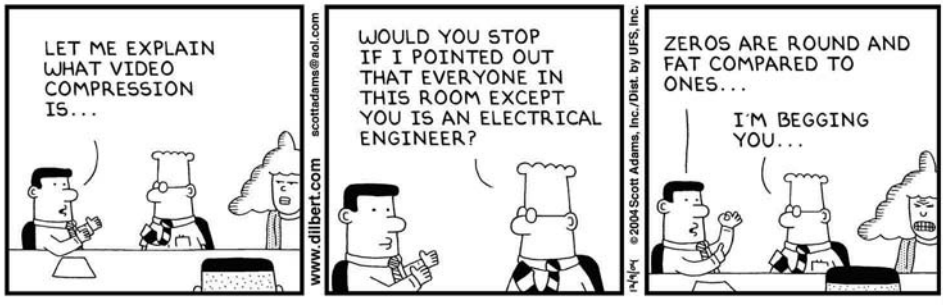

1530-LUMEN, 90W LIGHT ENGINE IS INCANDESCENT-EQUIVALENT

08.12.10

Light-engine manufacturer OptoElectronix recently introduced the 1530-lumen, 90W-incandescent-equivalent ULE5000 LightDisc light engine. In a circular form factor, the engine works well in both ceiling fans and wall mounts. It accepts 120V ac, has an 8.1-in. diameter with an integrated heat sink, and is compatible with TRIAC (triode-ac)-switch dimmers.

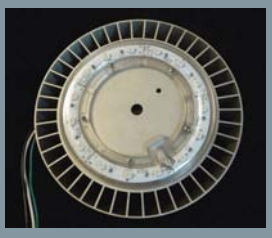

The LightDisc light engine includes an aluminum heat sink, with high-brightness LEDs faintly visible on the white substrate ring.

The engine touts a 50,000-hour life and is available in warm- or cool-white-light versions. According to the company, the device is three times more efficient than incandescent or halogen devices, and it consumes only 28W of power. It also features dynamic-thermal-management-control circuitry. The LightDisc LED light engine comes with a five-year or 35,000-hour warranty and sells for $129.20 (1000) each.

—by Margery Conner

▷OptoElectronix, www.optoelectronix.com.

The Power Behind the Green

High-efficiency, high-density power conversion

Vicor's powertrain and power management technologies enable extreme computing performance and efficiency.

Three of the top 10 supercomputers in the TOP500 list (TFLOPS) and twelve of the top 20 Green500 list (MFLOPS/Watt) are powered by Vicor.

Data centers using Vicor technology save $600k in electricity, 2,700 tons of CO_2, and 2,500 barrels of oil per year.

System features include:

- High-efficiency, high-voltage power distribution
- Direct conversion to processors, memory and drives
- Power conversion > 1000 W/in^3
- Highly efficient, very small size fault protection solutions for redundant power systems

Call 800-735-6200 to learn more about the power of Vicor technology.

V·I CHIP FACTORIZED POWER

vicorpower.com/datacenter/edn

VICOR

Novellus announces product suite for advanced WLP

Novellus Systems is addressing the challenges of 3-D chip integration with the introduction of new models of the company's Vector PECVD (plasma-enhanced chemically evaporated-deposition), Inova PVD (physical-vapor-deposition), and GxT photoresist-strip systems. The company also introduced the Sabre 3D electroplating system for the advanced-WLP (wafer-level-packaging) market.

The new PECVD, PVD, and photoresist-strip systems leverage Novellus' interconnect technology to meet the technical challenges of 3-D-integration technologies, such as TSVs (through-silicon vias) and WLP stacking schemes. TSVs allow manufacturers to vertically stack multiple chips and interconnect them across short distances with cross-device copper vias, leading to significant performance gains. WLP technologies, such as microbumps, pillars, and copper RDLs (redistribution layers), enable increasing I/O counts and decreasing pitch requirements. All of these 3-D-integration schemes are driving new metal and dielectric-interconnect materials and applications.

The Sabre 3D system incorporates Novellus' Sabre Electrofill technology and includes new patented technologies that enhance void-free filling, reduce copper overburden, and improve fill uniformity at high throughputs.

> WLP technologies enable increasing I/O counts and decreasing pitch requirements, driving new metal and dielectric-interconnect applications.

The Sabre 3D system incorporates Novellus' Sabre Electrofill technology and includes new patented technologies that enhance void-free filling, reduce copper overburden, and improve fill uniformity at high throughputs.

You can configure Sabre 3D's modular architecture with multiple plating and pretreatment or post-treatment cells for a variety of packaging applications, including TSV, pillar, RDL, underbump metallization, and eutectic and lead-free microbumping, using materials such as copper, tin, nickel, and tin silver.

The Inova PVD system uses Novellus' patented HCM (hollow-cathode-magnetron) sputtering source with the company's Ionflo technology to provide enhanced copper-sidewall coverage and low defects in high-aspect-ratio TSVs. The ion-induced copper flow process enables designers to achieve void-free fill using a much thinner seed layer than competitive PVD approaches. The new Vector system uses the same patented MSSD (multistation-sequential-deposition) technology that the company incorporated into previous-generation Vector PECVD tools.

The MSSD architecture works with the new PECVD processes to enable the system to deposit low-cost, low-temperature films that are compatible with bonded substrates. The system also enables the deposition of high-quality dielectric liners for TSV structures in both via-middle and via-last integration schemes. The new photoresist-strip system can quickly remove 20- to 100-micron-thick photoresists in the manufacture of RDLs and pillars and achieve residue-free strip and clean of high-aspect-ratio TSVs.

—by Rick Nelson

▷**Novellus**, www.novellus.com.

Virginia Tech launches smart-grid online-information portal

Virginia Polytechnic Institute and State University (Blacksburg, VA, www.vt.edu) recently announced that it has posted a beta version of its SGIC (Smart Grid Information Clearinghouse) Web portal, which will provide information on smart-grid technologies, standards, and rules and regulations.

Although the American Recovery and Reinvestment Act of 2009 (www.recovery.gov) primarily funds the smart grid, people are often unclear about its requirements, standards, and strategy. The Virginia Tech site, which will offer industry-use cases and case studies, public awareness and education information, and work-force-training opportunities, facilitates direct sharing and dissemination of smart-grid information among various stakeholders on knowledge gained, lessons learned, and best practices.

Virginia Tech describes the SGIC as a decision-support tool for both state and federal regulators in their deliberations for rule making and evaluating the impact of their investments in smart-grid technologies and software.

The SGIC came about through a $1.25 million, five-year contract that the US Department of Energy (www.energy.gov) awarded to the Virginia Tech ARI (Advanced Research Institute, www.ari.vt.edu) in the National Capital Region in October 2009. Virginia Tech developed the portal with content assistance from the IEEE (www.ieee.org) and EnerNex Corp (www.enernex.com).

> SGIC offers information for consumers who are just learning about the smart grid as well as industry professionals.

According to ARI Director Saifur Rahman, the Joseph R Loring professor of engineering at Virginia Tech and principal investigator for the SGIC portal, the SGIC offers information for consumers who are just learning about the smart grid, case studies and standards for industry professionals who are implementing new technologies, and in-depth technical specifications from those who have long been developing smart-grid technology. Over the next few months, he hopes many from these groups will offer comments and suggestions so that the official portal launch will appeal to a range of people. You can submit comments on the SGIC beta site at www.sgclearinghouse.org/?q=node/1191.

—by Suzanne Deffree

▷**Smart Grid Information Clearinghouse**, www.sgclearinghouse.org.

New light-bulb labels emphasize lumens, not watts

The days of specifying a light bulb's output in watts will soon pass, starting in mid-2011. The FTC (Federal Trade Commission) has announced a new Lighting Facts label

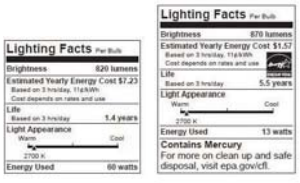

Light-bulb packages will soon carry one or both of these labels, listing light output in lumens rather than watts (left). CFL packages must also have a back label, disclosing the bulbs' mercury content (right).

for incandescent bulbs, which most manufacturers are phasing out over the next five years; CFLs (compact fluorescent lights); and LED-based lights (http://ftc.gov/opa/2010/06/lightbulbs.shtm).

When incandescent bulbs dominated light-bulb technology, describing a bulb's light output in watts made sense because the power the bulb used translated consistently to light output. The new labels may penalize higher-efficacy, lower-power-consumption CFLs and LEDs, however: A prospective buyer may think that a higher wattage rating means better light output.

The FTC requires manufacturers to place the new label on the front of lighting packages. The labels will list the bulb's brightness in lumens rather than watts. Packages for bulbs containing mercury, such as CFLs, must also include a back label indicating a Web site (www.epa.gov/cfl) that counsels consumers on the proper cleanup of broken bulbs and their safe disposal.

In addition to migrating lighting specification to a more relevant lumens, the FTC also mandates that the label must show the yearly energy cost, rather than the "yearly energy savings" that you see on many

> The new labels may penalize high-efficacy, low-power CFLs and LEDs: A prospective buyer may think that a higher wattage rating means better light output.

currently selling LED lights. For example, the ecoSmart LED bulb that Home Depot sells lists energy-cost savings of as much as $155 over its lifetime but doesn't state which technology it is comparing itself with. The new labels let consumers compare light output and yearly power costs for multiple technologies.

—by Margery Conner

▷**Federal Trade Commission**, www.ftc.gov.

Cree, Philips sign licensing agreement

Cree and Philips recently announced a cross-licensing agreement that covers patents from both parties in the fields of blue-LED-chip technology; white LEDs and phosphors, including remote phosphors; control systems; LED luminaires and lamps; and LED backlighting of LCDs and patents in the Philips LED Luminaire Licensing Program. As Cree gets more involved in lighting end products, this agreement could be an indicator of where the companies think LEDs in luminaires are heading.

Note that these are luminaires, not the 55W replacement lamp that Philips has submitted for the US Department of Energy's L-Prize (Lighting Prize, www.lightingprize.org). The L-Prize is the first government-sponsored technology competition to spur lighting manufacturers to develop high-quality, high-efficiency solid-state lighting products to replace common bulbs. **—by Margery Conner**

▷**Cree**, www.cree.com.

▷**Philips**, www.philips.com.

Online tool simplifies, speeds POL design

International Rectifier has announced a free online-design tool that enables electrical and thermal simulation for the company's SupIRBuck integrated POL (point-of-load) voltage regulators. The user-friendly, interactive, Web-based tool, available at http://mypower.irf.com/SupIRBuck, lets you quickly and easily select and simulate selected SupIRBuck products. Employing a designer's input and output parameters, the tool selects suitable devices for an application.

Once you enter basic requirements, the tool lets you capture schematics, create a reference design and an associated BOM (bill of materials), view waveforms, and quickly and easily perform complex thermal and application analysis to speed development time.

The tool is designed to assist a variety of designers—from power experts to digital specialists with little experience in analog design. The tool covers SupIR-Buck products with initial load currents as high as 12A, input voltages as high as 21V, and output voltages as low as 0.7V.

International Rectifier's online-design tool enables electrical and thermal simulation for the company's SupIRBuck family of integrated POL voltage regulators.

Support is available from customizable demo kits, reference designs, data sheets, and application notes. All of these materials are available at the company's Web site.

—by Paul Rako

▷**International Rectifier, www.irf.com.**

08.12.10

LINEAR TECHNOLOGY

DESIGN NOTES

2-Phase, Dual Output Synchronous Boost Converter Solves Thermal Problems in Harsh Environments – Design Note 481

Goran Perica

Introduction

Boost converters are regularly used in automotive and industrial applications to produce higher output voltages from lower input voltages. A simple boost converter using a Schottky boost diode (Figure 1) is often sufficient for low current applications. However, in high current or space-constrained applications, the power dissipated by the boost diode can be a problem especially in high ambient temperature environments. Heat sinks and fans may be needed to keep the circuit cool, resulting in high cost and complexity.

To solve this problem, the Schottky output rectifier can be replaced by a synchronous MOSFET rectifier (Figure 2). If MOSFETs with very low $R_{DS(ON)}$ are used, the power dissipation can be reduced to the point where no heat sinks or active cooling is required, thus reducing costs and saving space.

Advantages of Synchronous Rectification

Consider the power dissipation of the single output circuit in Figure 1. The output diode D1 carries 6.7A of RMS current to produce 3A of output current from a 5V input. At this current level, diode D1's voltage drop is 0.57V, resulting in 1.6W of power lost as heat. Dissipating 1.6W in an 85°C (or higher) automotive operating environment is not trivial. To keep the circuit cool, heat sinks, cooling fans and multilayer printed circuit boards must be used. This, of course, adds complexity, cost and size to an ostensibly simple boost converter.

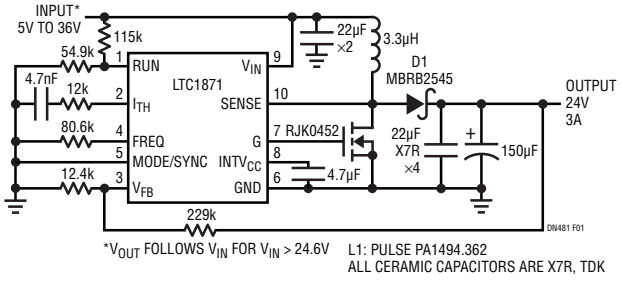

Figure 1. Although This Simple Circuit is Capable of 3A of Output Current, Beware of Power Dissipation in the Output Diode D1

A far better solution (featured in a dual output configuration) is shown in Figure 2, where a synchronous power MOSFET rectifier replaces the output diode. Under the same conditions, the voltage drop across output synchronous MOSFET Q2 is only 42mV or 7.4% of the voltage drop in the diode D1. The resulting power dissipation of 115mW in Q2 is relatively trivial. Another advantage of using a MOSFET as the output rectifier is the elimination of leakage current, about 10mA in the case of the MBR2545 diode—an additional 240mW of power dissipation in the application of Figure 1.

Dual Output Automotive Boost Converter

Figure 2 illustrates a typical automotive boost application with a 5V to 36V input voltage range. Here, the converter produces a 12V output for generic automotive loads such as entertainment systems, and a 24V output for circuits such as high power audio amplifiers. The two outputs are completely independent and can be controlled separately.

Because the circuit in Figure 2 is a boost converter, the output voltage can be regulated only for input voltages that are lower than the output voltage. The output voltage regulation versus input voltage is shown in Figure 3. When the input voltage is higher than the preset output voltage, synchronous MOSFETs Q2 and Q4 are turned continuously ON and boost MOSFETs are not switching. This feature allows the converter to be used in applications that require boosting only during load transients such as cold-cranking of a car engine. In this case, the LTC®3788 circuit's input voltage could be as low as 2.5V.

The efficiency of this converter (Figure 4) is high enough that it can be built entirely with surface mount components, requiring no heat sinks. A multilayer PCB with large copper area may be sufficient to dissipate the small amount of heat resulting from the MOSFETs' DC resistance, even at high ambient temperatures.

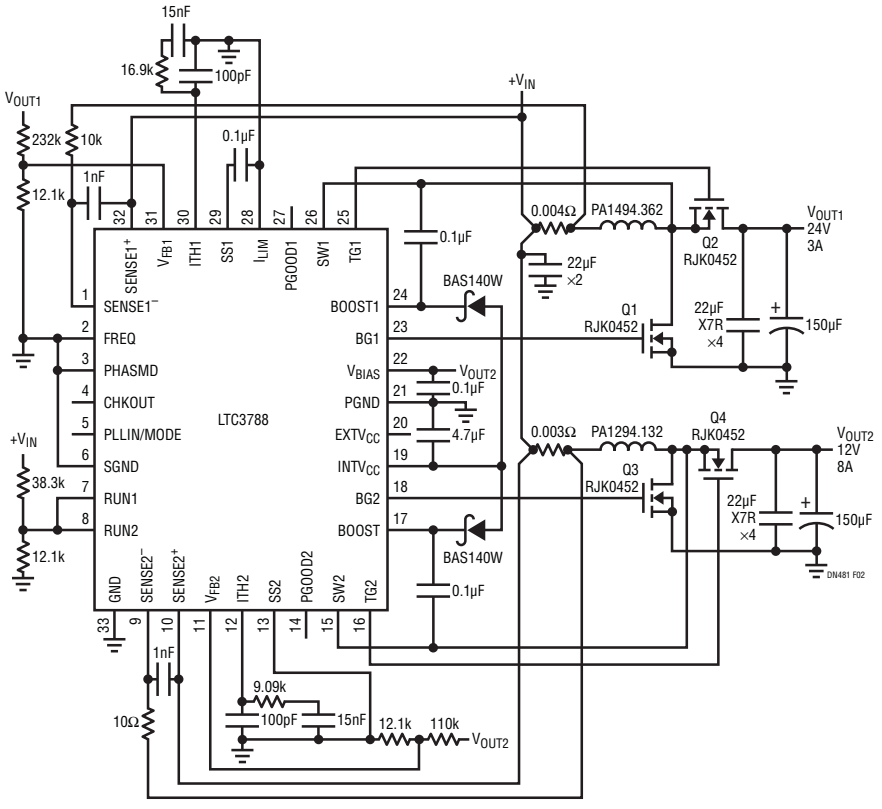

Figure 2. The LTC3788 Converter is Over 95% Efficient Even Under Worst-Case Conditions. When $V_{IN} > V_{OUT(SET)}$, Efficiency Approaches 100% as Shown in Figure 4

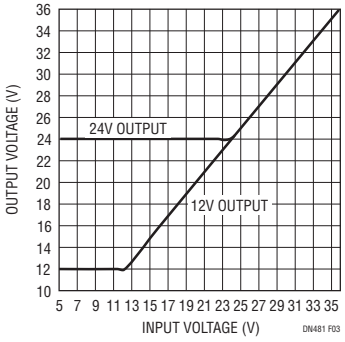

Figure 3. The Output Voltage Follows the Input Voltage when $V_{IN} > V_{OUT(SET)}$

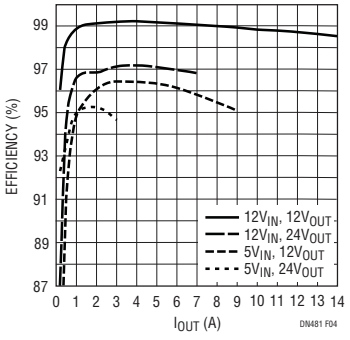

Figure 4. The Converter in Figure 2 Peaks at 95% Efficiency when Operating from a 5V Input

If higher output currents are required, or if lower output ripple voltage is desired, the two LTC3788 channels can be combined for a single current-shared output. Simply connect the two outputs and short the respective FB, ITH, SS and RUN pins. Because the two channels operate out of phase, output ripple currents are greatly reduced—nearly canceling out at 50% duty cycle. Thus, smaller output capacitors can be used with lower output ripple currents and voltages.

Conclusion

The LTC3788 dual synchronous boost controller is a versatile and efficient solution for demanding automotive and industrial applications. By minimizing power losses in the output rectifier, this converter can be designed in a very small footprint and operate safely at elevated ambient temperatures.

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 3788

Linear Technology Corporation
1630 McCarthy Blvd., Milpitas, CA 95035-7417
(408) 432-1900 • FAX: (408) 434-0507 • www.linear.com

dn481f LT/AP 0810 226K • PRINTED IN THE USA

BY HOWARD JOHNSON, PhD

Water hammer

For the second time in as many years, in the dead of night, a water buffalo lumbered into the village square and promptly fell into the community well. Ernie, the village engineer, argued that the citizens of the village should stop using the well and instead install a system of water pipes from the artesian spring on the mountainside behind the village.

He embellished his plan with a motor timer, a large electric snap-action water valve, and a flow regulator. The system would, he said, provide a reasonable, continuous flow of clean, potable water during the daytime. The excess flow could irrigate a system of flower beds, replacing those that the water buffalo's brief but deadly rampage had destroyed. At night, the valve would shut, removing from the central village square the nuisance of fresh water available to the local livestock.

The villagers approved. They raised a fund and launched Ernie on a journey to the nearest city to procure the necessary supplies. Ernie felt lucky to have this position. Tucked away high in the mountains of Bhutan, he could devote his engineering skill to the cause of social progress. He felt needed. The villagers built him a hut when he arrived. Here, he could forget all the unpleasantness of his last assignment.

When he returned from the city, Ernie labored for weeks burying 2500 feet of 4-in. steel irrigation pipe leading from the spring to the village far below.

He installed the valve and controller at the receiving end of the pipe. When the electric valve finally snapped open, water immediately gushed from the pipe at a rate of 30 gallons per minute, just as he had planned.

The villagers shouted and tossed their hats in the air. They plugged the old well with sand. They threw a feast. Later that night, as Ernie sat basking in the glow of his recent success, he heard the electric valve click shut. The pipe exploded, and an unstoppable mass of water, the whole unregulated flow of a 4-in. pipe, flooded the village.

Ernie failed to take into account the momentum of the enormous column of water moving inexorably toward the village; 2500 feet of 4-in. pipe holds about 13,000 lbs of water. That much water does not just suddenly stop moving because someone closed a snap-action valve.

Currents flowing in a long wire possess inductance, a property much like the momentum of flowing water.

You must slowly shut down these types of systems—over a much longer time than the round-trip delay of pressure waves traveling end to end throughout the system. At a propagation velocity of almost 5000 feet/sec, the round-trip propagation delay of pressure waves traveling there and back through Ernie's column of water is about 1 second. Slowly closing the valve over a period of, say, 10 seconds would reduce by a factor of 10 the peak amount of pressure needed to stop the movement of all that water.

Electrical circuits obey a similar principle. Currents flowing in a long wire possess inductance, a property much like the momentum of flowing water. Suddenly arrest the flow of current at the end of a long wire, and you will see a brief but intense inductive-voltage spike. If the spike is sufficiently large, it can destroy your circuit just as easily as Ernie's plumbing blew up his chance at a successful second career. EDN

Howard Johnson, PhD, of Signal Consulting, frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. Visit his Web site at www.sigcon.com, or e-mail him at howie03@sigcon.com.

BY PALLAB CHATTERJEE, CONTRIBUTING TECHNICAL EDITOR

Nano drives improve displays

Display technology is currently realizing the benefits of nanotechnology in lighting support for the displays and the display construction itself. One of the new display technologies is the Mirasol display (www.mirasoldisplays.com) from Qualcomm (www.qualcomm.com). This MEMS (microelectromechanical-system)-on-glass device targets low-power, daylight-readable color displays for portable-system applications.

Most LCD devices operating at low power, such as with mobile phones and tablet PCs, have issues with color representation. In varying light, the color accuracy of the display changes, altering the viewer's perception of the image. The Mirasol display attempts to overcome these issues.

The display is a front-reflective display rather than a traditional backlit display. The properties of nanoscale materials combine with advanced MEMS-processing techniques, allowing the display to mimic naturally occurring phenomena. The display works by creating a color from an interference pattern on the reflected light that hits the top of the display. This process is the same one that makes a butterfly's wing shimmer and display different colors.

The display uses red, green, and blue subpixels to create a single pixel measuring less than 1 micron. Each of the colored subpixels comprises cells, which form an array of devices. An applied voltage then switches these devices to "collapse" the MEMS device and turn off the reflective element (**Figure 1**).

The Mirasol display is just one example of the use of nanotechnology in a diverse range of applications. The area of nanomaterials—compounds, elements, and polymers that behave differently or exhibit new and unique properties at the nanoscale—is bringing many new support applications to electronics.

In the case of products from Cima NanoTech (www.cimananotech.com), the key nanomaterial is nanocrystalline silver. With its high conductive characteristics, silver has been in use as a standard material for more than 100 years. Cima formulates the silver and silver/copper into nanoparticle form and places these particles in suspensions and emulsions that designers can apply to various surfaces. Using self-assembly techniques, these nanoparticles form microscopic conductive networks that are transparent but that have properties similar to those of traditional opaque materials. Designers can apply the nanoparticle layer as a transparent film with an inexpensive, wet-rolling process in a large format, so it can become a low-cost, effective EMI (electromagnetic-interference) filter for plasma and other large displays. Designers can also pattern the material to form large-format resistive-touch-sensor overlays for LCD screens. In a new and growing application, the silver-and-copper particle solution can find use in conductive inks. Designers are able to print these inks on a surface using time-proven ink-jet technology.

This ability to print a conductor is drastically changing the RF- and PCB (printed-circuit-board)-prototyping flows by allowing, for example, the creation of one-time unique antenna designs that designers print and test on paper. Designers can similarly print a few fine-pitch prototypes on paper or mylar to test a design. The technique currently works in single-layer-system applications. However, multilayer-system applications are in progress.

Quantum dots are nanoparticles that have tunable optical characteristics. Companies including Nanosys (www.nanosysinc.com) and Zymera (www.zymera.com) have turned these properties into commercial products. Nanosys has a series of quantum-dot lighting techniques that color-correct the wavelength of backlit LEDs in cell phones. Zymera is using the dots as bioluminescent particle tags for drugs to track the absorption and location of certain medications in the treatment of cancer, for example. EDN

Contact me at pallabc@siliconmap.net.

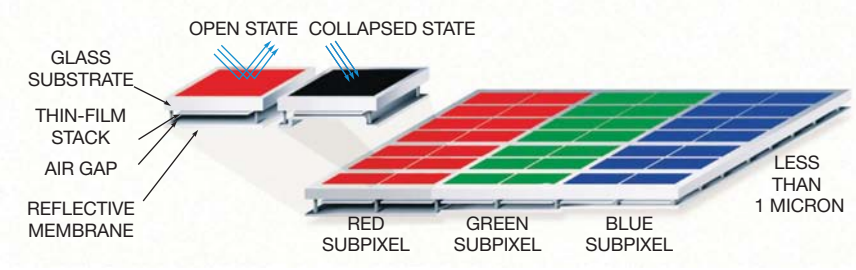

Figure 1 The Mirasol display uses MEMS on glass to create colors using interference patterns. The red subpixel has an open state and a collapsed state. The red-, green-, and blue-pixel design using MEMS receives light modulation from applied voltage, color selection from constructive interference, and memory from electromechanical behavior.

The Trade Shows that Connect You with the Electronics Manufacturing Industry

Electronics® IPC® Midwest

September 28–30, 2010
Donald E. Stephens Convention Center
Rosemont (Chicago), IL
ElectronicsMidwest.com

Electronics® WEST

February 8–10, 2011
Anaheim Convention Center
Anaheim, CA
ElectronicsWestShow.com

Electronics® NEW ENGLAND

April 6–7, 2011
Boston Convention & Exhibition Center
Boston, MA
Electronics-NE.com

Find the suppliers, tools, and services you need to make your product, process, and business more efficient, cost-effective, and profitable.

For information on attending or exhibiting, please call 310/445-4200 or visit CanonTradeShows.com

15512_AS_EL10

Produced and managed by Canon Communications LLC • 11444 W. Olympic Blvd. • Los Angeles, CA 90064-1549 • Tel: 310/445-4200

BY JIM WILLIAMS • LINEAR TECHNOLOGY

MEASURING WIDEBAND-AMPLIFIER SETTLING TIME

A NOVEL CIRCUIT LETS YOU MEASURE OUTPUT SETTLING TO 0.1% IN 2 nSEC.

You use wideband amplifiers in instrumentation, waveform-synthesis, data-acquisition, and feedback-control systems. To ensure a robust design for these systems, you must verify precision operation at high speeds. This requirement presents a difficult measurement challenge. Wideband operational amplifiers feature dc precision of 0.2-mV offset voltage with gain-bandwidth products of 400 MHz and slew rates of 2500V/µsec (**Reference 1**). IC designers face a trade-off between fast slew rates and short ring times. Fast-slewing amplifiers generally have extended ring times. This combination complicates your amplifier choice and the frequency compensation you use (see **sidebar** "Practical considerations for amplifier compensation"). Additionally, the architecture of very fast amplifiers usually dictates trade-offs, which degrade dc-error terms.

SETTLING TIME DEFINED

It is relatively easy for you to verify amplifier dc specifications. Literature defines the measurement techniques you use. You need to use more sophisticated approaches to produce reliable ac specifications. Measuring anything at any speed requires care. Dynamic measurements are particularly challenging, and amplifier settling time is difficult to determine (**references 2** through **7**).

Settling time is the elapsed time from an input application until the output arrives at and remains within a specified error band around the final value. Amplifier manufacturers usually specify it over a full-scale transition.

Settling time has three distinct components (**Figure 1**). The delay time is small and almost entirely due to the amplifier's propagation delay. No output movement occurs during this in-

terval. During slew time, the amplifier moves at its highest speed toward the final value. Ring time defines the region during which the amplifier recovers from slewing and ceases movement within some defined error band. Measuring settling times of nanoseconds requires a careful approach and experimental technique.

The traditional way for measuring settling time is with a circuit that uses the false-sum-node technique (**Figure 2**). The resistors and amplifier form a bridge network. The amplifier output steps to the input voltage when you drive the input, assuming that the circuit is using ideal resistors. During the slew period, the diodes bound the settle node, limiting the voltage excursion. When settling occurs, the oscilloscope's probe voltage should be 0V. The resistor divider's attenuation causes the probe's output to be one-half of the settled voltage.

In theory, this circuit should allow you to observe fast settling to small amplitudes. In practice, you cannot rely on it to produce useful measurements. The circuit has several flaws, including a requirement for the input pulse to have a flat top within the required measurement limits. Typically, you are interested in settling of less than 5 mV for a 5V step. No general-purpose pulse generator holds the output amplitude and noise within these limits. You cannot distinguish between generator-caused aberrations and amplifier-related ones.

The oscilloscope connection also presents problems. As probe capacitance rises, the ac loading of the resistor junction influences the observed settling waveforms. The excessive input capacitance of 1× probes makes them unsuitable for this measurement. A 10× probe's attenuation sacrifices oscilloscope gain, yet its 10-pF input capacitance still introduces a significant lag at nanosecond speeds. If you use an active 1×, 1-pF FET (field-effect-transistor) probe, it largely alleviates this problem, but a more serious issue remains.

You use clamp diodes at the settle node to reduce the voltage swing during amplifier slewing. This approach is intended to prevent the circuit from overdriving the oscilloscope input. Unfortunately, the 400-mV drop across the Schottky diodes means that the oscilloscope will undergo an unacceptable overload (**Reference 8**). Oscilloscopes' overdrive-recovery characteristics vary widely among models and brands, and manufacturers typically do not specify it. At 0.1% resolution, the oscilloscope typically undergoes a 10-times overdrive at 10 mV/division, making the desired 2.5-mV baseline unattainable.

AT A GLANCE

- It is difficult to measure the settling time of fast amplifiers.
- One method of measurement is to use an old analog sampling oscilloscope.
- This circuit uses a current-steering bridge and an analog multiplier IC.
- The two methods correlate well.
- The measurement limits are 2 mV and 2 nsec.

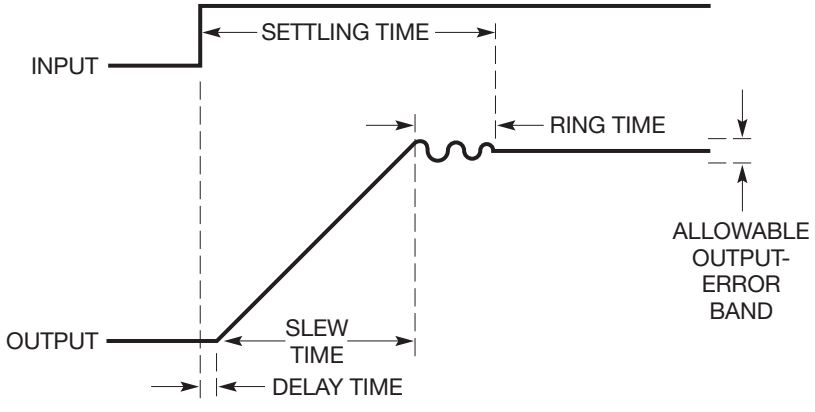

Figure 1 The components of settling time include delay, slew, and ring times. Using fast amplifiers reduces slew time, although longer ring time usually results.

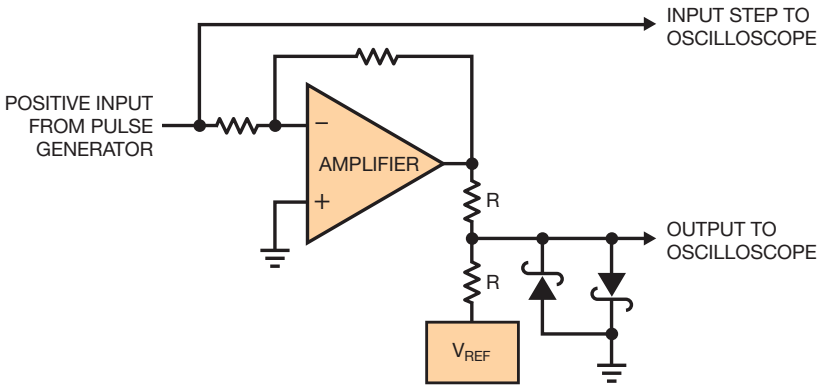

Figure 2 A popular summing scheme for settling-time measurement provides misleading results with fast amplifiers.

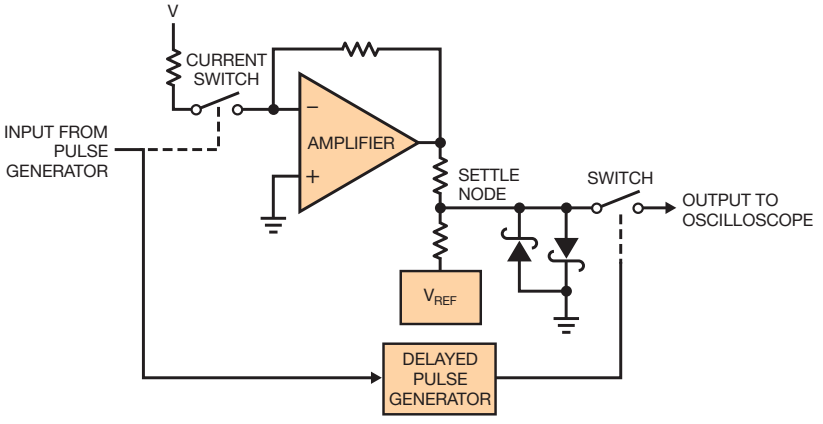

Figure 3 The proposed circuit's conceptual arrangement is insensitive to pulse-generator aberrations and eliminates oscilloscope overdrive. The input switch gates a current step to the amplifier under test. A delayed-pulse generator controls a second switch, which prevents the oscilloscope from monitoring the settle node until settling is nearly complete.

With this arrangement, the measurement becomes hopeless at nanosecond speeds.

Thus, your measuring wideband amplifier settling time requires an oscilloscope that is somehow immune to overdrive as well as a flat-top pulse generator. The only oscilloscope technology that offers inherent overdrive immunity is the classic analog sampling oscilloscope. Do not confuse these scopes with modern digital sampling oscilloscopes that have overdrive restrictions (**Reference 8**). Several documents explain the operation of classic sampling oscilloscopes (**references 9** through **13**). Although you can buy these instruments used, their manufacturers no longer make them. You can, however, construct a circuit that borrows the overload advantages of classic analog-sampling-oscilloscope technology. You can also endow the circuit with features for measuring nanosecond settling times.

You can avoid the flat-top-pulse-generator requirement by switching current rather than voltage. It is easier to gate a quickly settling current into the amplifier's summing node than to control a voltage. This approach eases the input pulse generator's job, although it still must have a rise time of approximately 1 nsec to avoid measurement errors.

PRACTICAL MEASUREMENT

A circuit that can measure wideband-amplifier settling time shares attributes

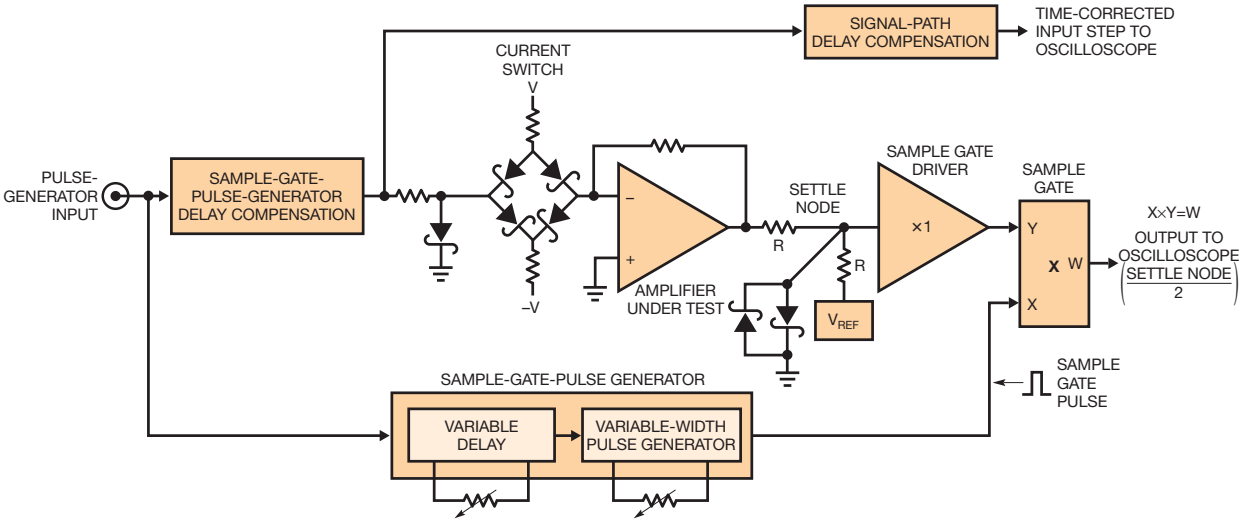

Figure 4 The diode bridge cleanly switches input current to the amplifier. A multiplier-based sampling switch eliminates the signal's presettling excursion. You compensate the input-step time reference and sample-gate-pulse generator for any circuit delays.

with the classic method, although some new features appear (**Figure 3**). The oscilloscope connects to the settle point by a switch. You determine the switch's state by triggering a delaying pulse generator from the input pulse. You arrange the delayed pulse generator's timing so that the switch does not close until settling is nearly complete. In this way, you sample the incoming waveform in both time and amplitude. No off-screen activity occurs on the oscilloscope; hence, you never subject the oscilloscope to overdrive.

You control the switch at the amplifier's summing junction with the input pulse. This switch gates current to the amplifier through a voltage-driven resistor. This approach eliminates the requirement for a flat-top pulse generator, although the switch must be fast and devoid of drive artifacts.

THE SAMPLE-GATE MULTIPLIER IC MUST PASS WIDEBAND-SIGNAL-PATH INFORMATION WITHOUT INTRODUCING ALIEN COMPONENTS.

For more detail, you split the delayed pulse generator into a delay block and a pulse generator, which you can vary independently (**Figure 4**). The input step to the oscilloscope runs through a section that compensates for the propagation delay of the settling time's measurement path. Similarly, another delay compensates the sample gate's pulse-generator propagation delay. This delay causes a phase-advanced version of the pulse that triggers the amplifier under test to drive the sample gate's pulse generator. This approach improves minimum measurable settling time by making irrelevant the sample gate's pulse-generator propagation delay.

The most striking new aspects of **Figure 4**'s circuit are the diode bridge switch and the multiplier IC. The diode bridge's balance combines with matched, low-capacitance Schottky diodes and high-speed drive to yield clean switching. The bridge quickly switches current into the amplifier's summing point, with settling time within 1 nsec. The diode clamp to ground prevents excessive bridge-drive swings and ensures that nonideal input-pulse characteristics are nearly irrelevant.

The sample-gate multiplier IC has stringent requirements. It must faithfully pass wideband-signal-path information without introducing alien components, particularly those deriving from the switch-command channel that provides the sample-gate pulse. Conventional choices for the sample-gate switch would include FETs or a sampling diode bridge. But FETs' parasitic gate-to-channel capacitances would result in large gate-drive-originated feedthrough into the signal path. For almost all FETs, this feedthrough is many times larger than the signal you are observing and would induce oscilloscope overload and obviate the switch's purpose. The diode bridge is better; its small parasitic capacitances tend to cancel, and its symmetrical, differential structure results in low feedthrough. Practically, however, the bridge requires dc and ac trims and complex drive and support circuitry (**references 3**, **4**, **7**, and **14**).

To avoid these problems, the sample-gate multiplier IC functions as a wideband high-resolution switch with low feedthrough. The great advantage of this approach is that you can maintain the switch-control channel inband. You hold the transition rate within the multiplier IC's 250-MHz bandpass. The multiplier's wide bandwidth means that you always control the switch command's transition. There are no out-of-band responses, greatly reducing feedthrough and parasitic artifacts.

SETTLING-TIME CIRCUITRY

You let the input pulse switch the input bridge through a delay network of inverters, A, and a driver stage comprising similar inverters, C (**Figure 5**). The delay compensates the sample-gate pulse generator's delayed response. This step ensures that the sample gate's pulse can occur immediately after the end of the amplifier under test's slew time. You choose the delay range so that the sample-gate pulse can occur before the amplifier slews. This capability is unused in normal operation, although it guarantees that you will always be able to

capture the settling interval.

The C inverters form a noninverting driver stage you use to switch the diode bridge. You adjust various trims to optimize the driver output pulse shape (see **sidebar** "Settling-time circuit-trimming procedures" with the Web version of this article at www.edn.com/100812df). This approach provides a clean, fast impulse to the diode bridge. This high-fidelity pulse is devoid of undamped components. It prevents radiation and disruptive ground currents from degrading the measurement noise floor. The driver also activates the B inverters, which supply a time-corrected input step to the oscilloscope.

THE PULSE IS DEVOID OF UNDAMPED COMPONENTS AND PREVENTS RADIATION FROM DEGRADING THE MEASUREMENT NOISE FLOOR.

The driver's output pulse transitions through the 1N5712 diode clamp's forward-voltage potential in less than 1 nsec. This transition causes an essentially instantaneous switching of the diode bridge. The cleanly settling current into the amplifier under test's sum-

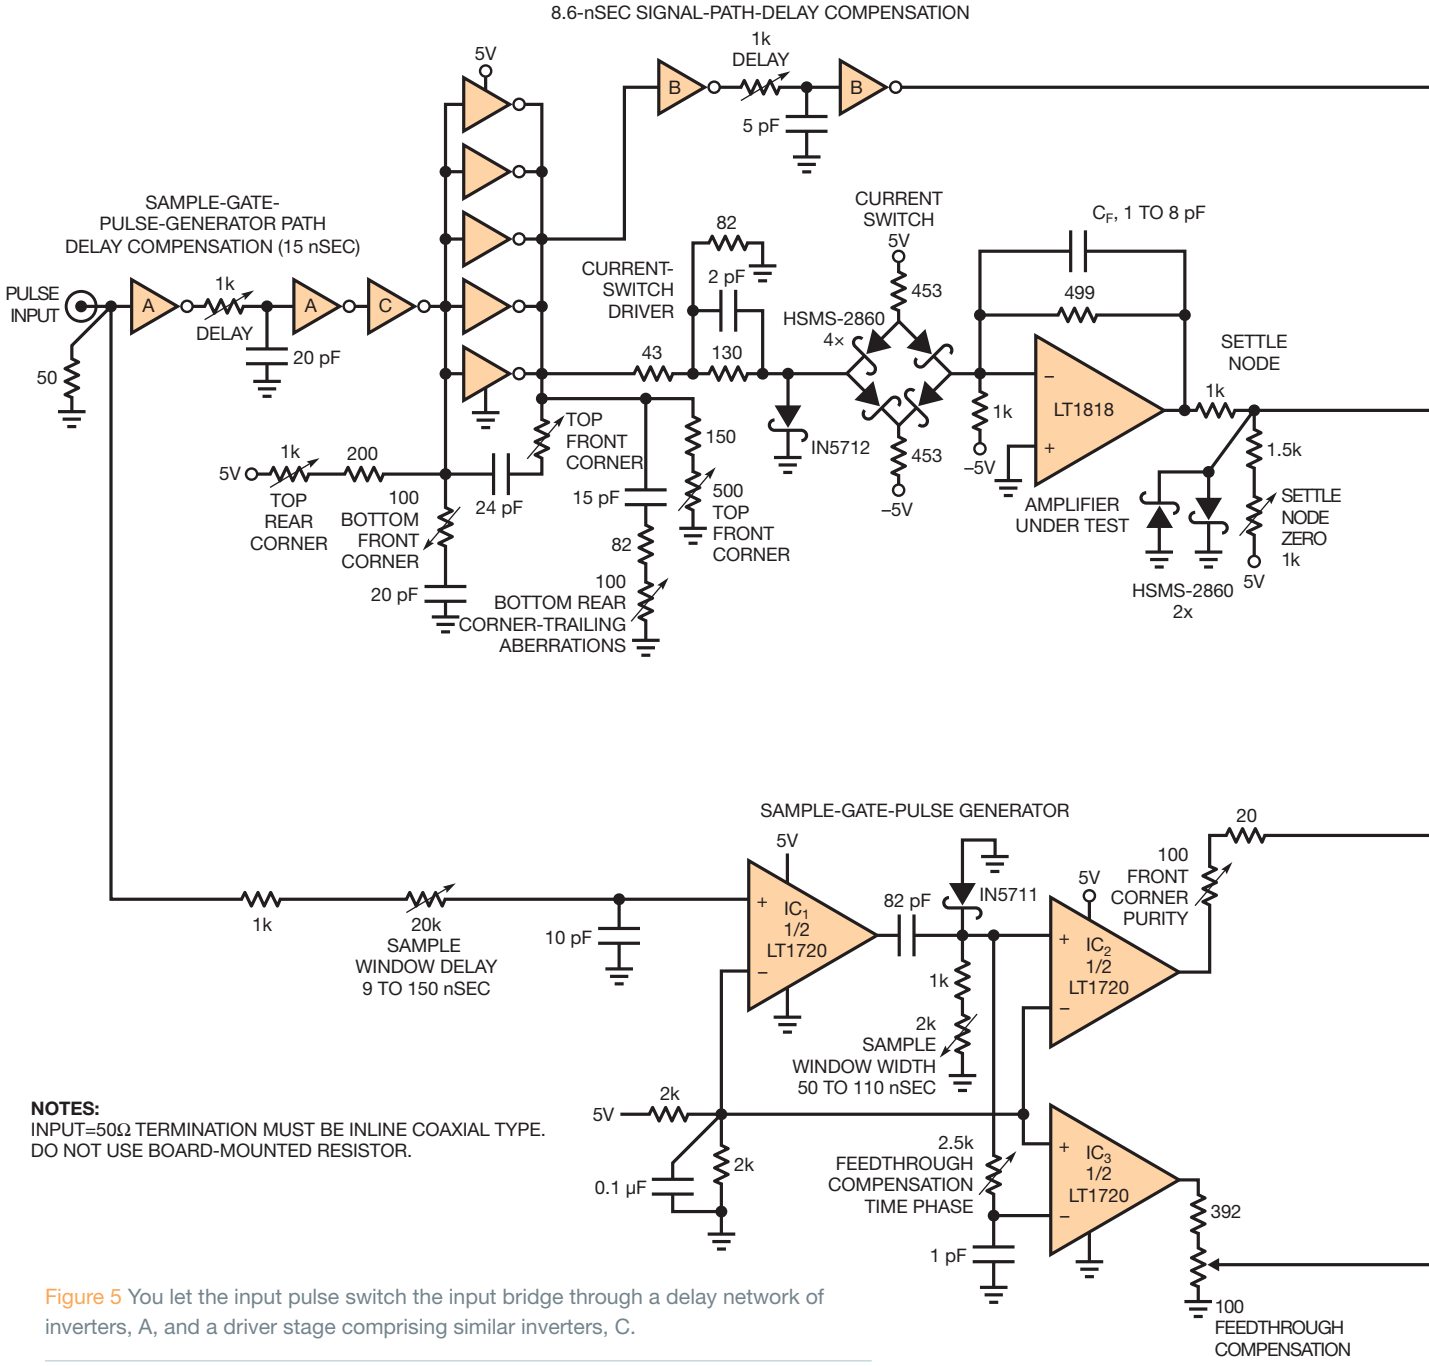

NOTES:
INPUT=50Ω TERMINATION MUST BE INLINE COAXIAL TYPE.
DO NOT USE BOARD-MOUNTED RESISTOR.

Figure 5 You let the input pulse switch the input bridge through a delay network of inverters, A, and a driver stage comprising similar inverters, C.

ming point causes a proportionate amplifier output movement. You set up a negative bias current at the amplifier's summing point with a 1-kΩ resistor pulled to −5V. That current combines with the input current step to produce a −2.5 to +2.5V amplifier output transition. You feed this amplifier's output to a voltage divider biased to 5V. You adjust the potentiometer to a nominal 500Ω so that when the amplifier under test transitions to −2.5V, the node clamped by the two Schottky diodes transitions to 0V. Buffer amplifier A_1 unloads this clamped settle node and provides the settling-time signal to the AD835 multiplier IC.

The other signal path to the multiplier IC uses a 20-kΩ potentiometer to set a delay time of the input pulse. This potentiometer feeds three comparators, and you use a 2-kΩ potentiometer to set the delayed pulse width. This step sets the sample gate's on-time. The Q_1 stage forms the sample gate's pulse into a clean, fast rise time. This technique furnishes pure, calibrated-amplitude, on/off switching instructions to the sample gate's multiplier IC. Appropriate setting of the sample gate's pulse delay means that the oscilloscope will not see any input until settling is nearly complete, eliminating oscilloscope overdrive. You adjust the sample window's pulse width so that you can observe all the remaining settling activity. In this way, the oscilloscope's output is reliable, and you can take meaningful data.

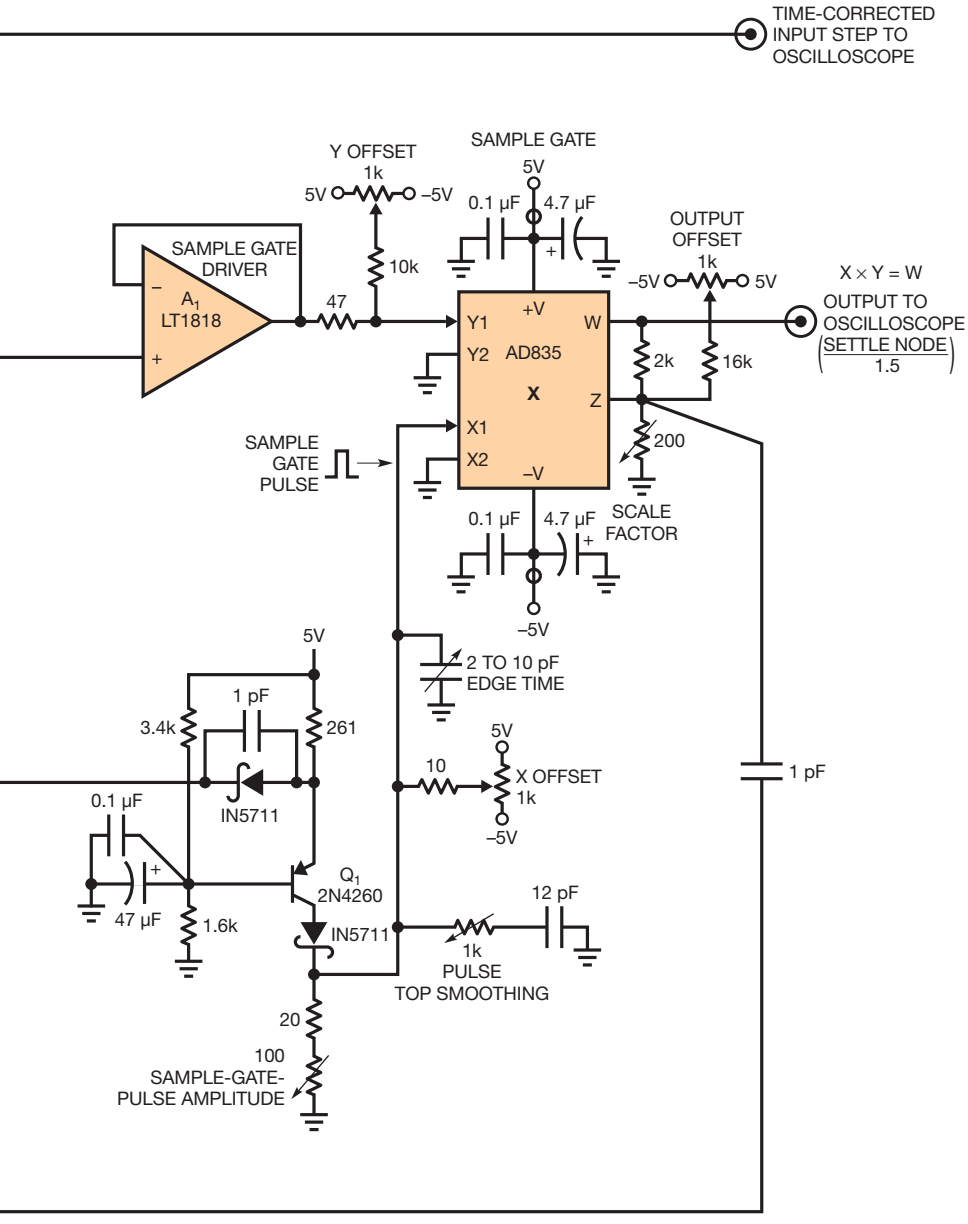

PERFORMANCE RESULTS

The circuit performs admirably (**Figure 6**). Trace A is the time-corrected input pulse, Trace B is the amplifier's output, Trace C is the sample gate's pulse, and Trace D is the settling-time output. When you interpret the waveform placement, note that Trace B appears time-skewed relative to time-corrected Trace A. This skew accounts for Trace B's false movement before Trace

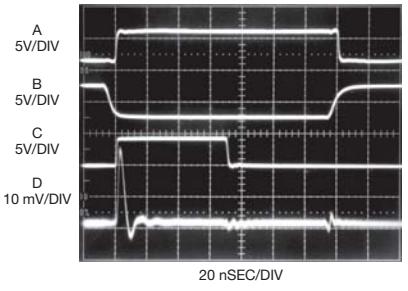

Figure 6 The settling-time circuit's waveforms include the time-corrected input pulse (Trace A), amplifier under test's output (Trace B), sample gate's trace (Trace C), and settling-time output (Trace D). You can vary the sample-gate window's delay and width. Trace B appears time-skewed relative to the time-corrected Trace A.

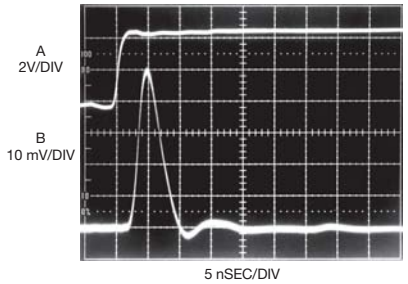

Figure 7 By expanding the vertical and horizontal scales, you can see the 9-nsec amplifier settling to within 5 mV (Trace B). Trace A is the time-corrected input step.

A's ascent. When the sample gate's pulse goes high, the sample gate switches cleanly. You can easily observe the last 20 mV of the amplifier's slewing. You can also see the entire ring time and the amplifier settling nicely to a final value. When the sample gate's pulse goes low, the sample gate switches off with only 2 mV of feedthrough. No off-screen activity occurs at any time, and you never subject the oscilloscope to overdrive.

You can adjust the vertical and horizontal scales of the oscilloscope to make the settling details more visible (**Figure 7**). You measure settling time from the onset of the time-corrected input pulse. Additionally, you calibrate the settling signal's amplitude with respect to the

PRACTICAL CONSIDERATIONS FOR AMPLIFIER COMPENSATION

There are a number of practical considerations when you compensate the amplifier to get the fastest settling time (see Figure 1 in the main text). Once you choose an amplifier, the only settling variable you can change is the ring time by changing the amplifier's compensation network. Because slew time is usually the dominant lag, it is tempting to select the fastest-slewing amplifier available. Unfortunately, fast-slewing amplifiers usually have extended ring times, which negate their speed advantage. If you damp out the ringing with large compensation capacitors, it results in protracted settling times.

The key to good settling times is to choose an amplifier with the right balance of slew rate and recovery characteristics and then properly compensate it. This is harder than it sounds because you can't predict amplifier settling time or extrapolate it from any combination of data-sheet specifications. You must measure settling time in the intended configuration.

A number of terms combine to influence settling time. They include amplifier slew rate and ac dynamics, layout capacitance, source resistance and capacitance, and the compensation capacitor. These terms interact in a complex manner, making predictions hazardous. Spice aficionados should take notice. If you replace the parasitic components with a purely resistive source, you still can't readily predict amplifier settling time. The parasitic-impedance terms make a difficult problem messier.

The only way to deal with the parasitic terms is to adjust the feedback compensation capacitor, C_F. When you properly adjust the value of C_F, it rolls off the amplifier's gain at the frequency that permits the best dynamic response. You achieve the best settling results when you select the capacitor to functionally compensate for all the terms (Figure A).

Trace A is the time-corrected input pulse, and Trace B is the amplifier's settle signal. The amplifier comes cleanly out of slew and settles to 5 mV in 9 nsec. The sample gate opens just after the second vertical division. The waveform signature is tight and nearly critically damped. When you use too large a feedback capacitor the settling is smooth, although overdamped (Figure B), giving you a 13-nsec penalty that results in a 22-nsec settling time. Eliminating the feedback capacitor results in a severely underdamped response with resultant excessive ring-time excursions (Figure C). Settling time goes out to 33 nsec. Using a feedback capacitor that is too small results in an underdamped response requiring

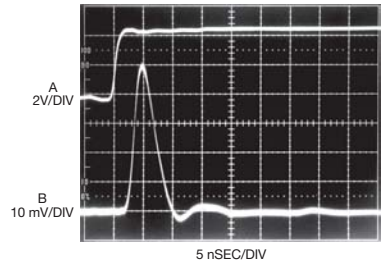

Figure A Optimizing the compensation capacitor permits a tight waveform signature, a nearly critically damped response, and a 9-nsec settling time. Trace A is the time-corrected input step. Trace B is the settle signal.

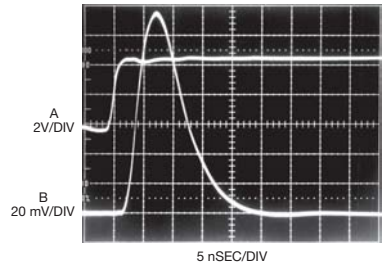

Figure B This overdamped response ensures freedom from ringing, even with component variations in production. The penalty is a 22-nsec settling time. Note the two-times vertical-scale change versus that of Figure A.

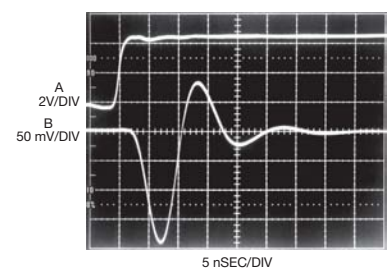

Figure C This severely underdamped response is due to the lack of a feedback capacitor. Note the five-times vertical-scale change versus that of Figure A. Settling time is 33 nsec.

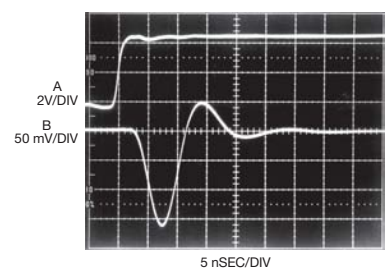

Figure D An underdamped response results from an undersized capacitor. Component-tolerance budgeting prevents this behavior. Note the five-times vertical-scale change versus that of Figure A. Settling time is 27 nsec.

amplifier, not the settle node. This approach eliminates ambiguity due to the settle node's resistor ratio. Trace A is the time-corrected input pulse, and Trace B is the settling output. You can easily observe the last 50 mV of slew. The amplifier settles within 5 mV, or 0.1%, in 9 nsec after you optimize the amplifier under test's feedback capacitor, C_F, (see **sidebar** "Practical considerations for amplifier compensation").

It is good practice to adjust the sampling window backward to the last 50 mV or so of amplifier slewing. This step allows you to observe the onset of ring time without encountering oscilloscope overdrive. The sampling-based approach provides this capability, and it is a powerful measurement tool. Slower amplifiers may require extended delay, sampling-window times, or both. You can use larger capacitor values in the delayed pulse-generator timing networks to meet these requirements.

TO ENSURE CONFIDENCE, MAKE THE SAME MEASUREMENT WITH AN ALTERNATIVE METHOD AND SEE WHETHER RESULTS AGREE.

VERIFYING RESULTS

The sampling-based settling-time circuit appears to be a useful measurement approach. A good way to ensure confidence is to make the same measurement with an alternative method and see whether results agree.

Classic sampling oscilloscopes are inherently immune to overdrive (**Reference 8**). You can use this feature and attempt a settling-time measurement directly at the clamped settle node (**Figure 8**). The circuit heavily overdrives a Tektronix type 661 with 4S1 vertical and 5T3 timing plug-ins, but the instrument is ostensibly immune to the insult (**Figure 9**). Trace A is the time-corrected input pulse, and

27 nsec to settle (Figure D). Note that figures B, C, and D require you to reduce the vertical scale to capture the nonoptimal responses.

When you trim the feedback capacitor for optimal response, the source, stray, amplifier, and compensation-capacitor tolerances are irrelevant. If you don't use individual trimming, you must consider these tolerances to determine the feedback capacitor's production value. Stray and source capacitance and output loading, as well as the feedback capacitor's value, affect ring time. The relationship is nonlinear, although some guidelines are possible. The stray and source terms can vary by ±10%, and the feedback capacitor is typically a ±5% component. These figures assume a resistive source. If the source has substantial parasitic capacitance, such as a photodiode or a DAC, this number can easily reach ±50%. Additionally, amplifier slew rate has a significant tolerance, which the data sheet states. To obtain a production-feedback-capacitor value, you determine the optimum value by individual trimming with the production-board layout. Remember that board-layout parasitic capacitance counts too. Then, factor in the worst-case percentage values for stray- and source-impedance terms, slew rate, and feedback-capacitor tolerance. Add this information to the trimmed capacitors' measured value to obtain the production value. This budgeting is perhaps unduly pessimistic. The errors should sum in an rms (root-mean-square) fashion, not a purely additive way.

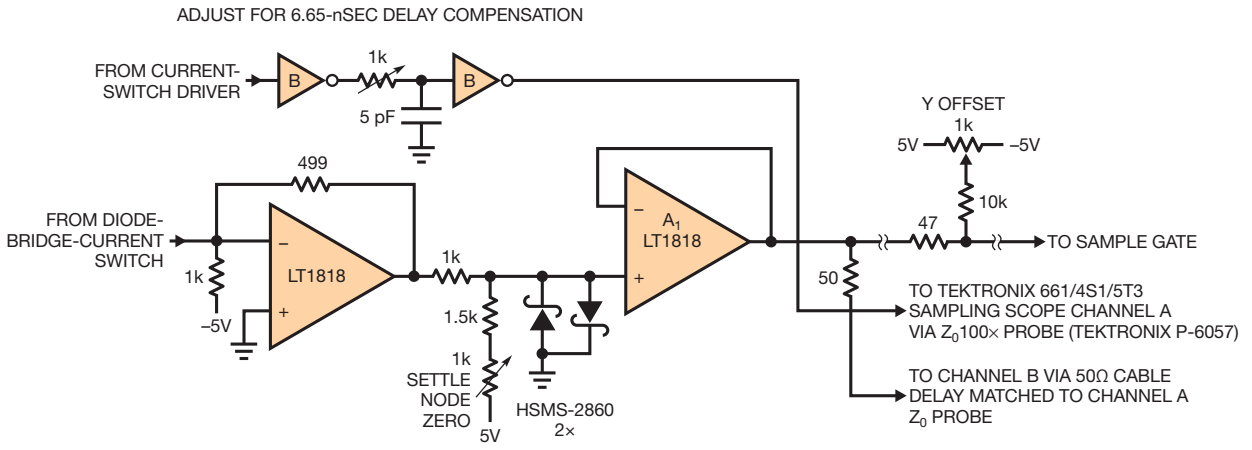

Figure 8 This alternative measurement method uses a 1-GHz Tektronix 661/4S1/5T3 sampling oscilloscope.

Trace B is the settle signal. Despite a brutal overdrive, the oscilloscope responds cleanly, giving a plausible settle signal.

You can compare the results visually (**figures 9** and **10**). Ideally, if both approaches represent good measurement technique and you properly construct them, the results should be identical. If this scenario occurs, the data produced by the two methods has a high probability of being valid.

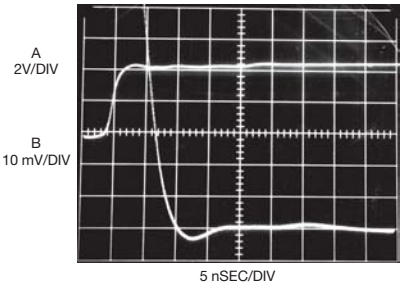

Figure 9 The measurement using a sampling oscilloscope shows a 9-nsec settling time and waveform profile, which are consistent with those of Figure 7.

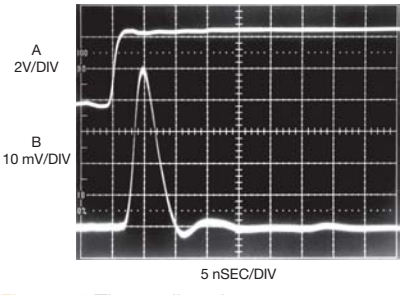

Figure 10 The settling-time measurement using the circuit of Figure 5 yields results that correlate well with those of Figure 9.

The two measurement methods do show nearly identical settling times and highly similar settling-waveform signatures. This agreement provides a high degree of credibility to the measured results. The noise floor and the signal feedthrough impose the 2-mV amplitude-resolution limit. The time resolution's limit is about 2-nsec to 5-mV settling. EDN

REFERENCES

1 "LT1818/LT1819 400MHz, 2500V/µs, 9mA Single/Dual Operational Amplifiers," Linear Technology Corp, 2002, http://bit.ly/cQ0xUe.

2 Williams, Jim, "1ppm Settling Time Measurement for a Monolithic 18-Bit DAC," Linear Technology Corp, Application Note 120, March 2010, http://bit.ly/aqLavq.

3 Williams, Jim, "30 Nanosecond Settling Time Measurement for a Precision Wideband Amplifier," Linear Technology Corp, Application Note 79, September 1999, http://bit.ly/9yApYn.

4 Williams, Jim, "Component and Measurement Advances Ensure 16-Bit DAC Settling Time," Linear Technology Corp, Application Note 74, July 1998, http://bit.ly/dePR4U.

5 Williams, Jim, "Measuring 16-Bit settling times: the art of timely accuracy," *EDN*, Nov 19, 1998, pg 159, http://bit.ly/c8xAS5.

6 Williams, Jim, "Methods for Measuring Op Amp Settling Time," Linear Technology Corp, Application Note 10, July 1985, http://bit.ly/bfnrU1.

7 Kayabasi, Cezmi, "Settling Time Measurement Techniques Achieving High Precision at High Speeds," Worcester Polytechnic Institute, 2005, http://bit.ly/cNtCp0.

8 Williams, Jim, "Precisely measure settling time to 1 ppm," *EDN*, March 4, 2010, pg 20, http://bit.ly/b9ITLG.

9 Korn, Granio A, and Theresa M Korn, *Electronic Analog and Hybrid Computers*, "Diode Switches," pg 223, McGraw-Hill, 1964.

10 Carlson, R, "A Versatile New DC-500 MHz Oscilloscope with High Sensitivity and Dual Channel Display," *Hewlett-Packard Journal*, January 1960.

11 Tektronix Inc, *Type 1S1 Sampling Plug-In Operating and Service Manual*, Tektronix Inc, 1965.

12 Mulvey, John, *Sampling Oscilloscope Circuits*, Tektronix Inc, Concept Series, 1970.

13 "Sampling Notes," Tektronix Inc, 1964, http://bit.ly/9agxbe.

14 Pease, Robert A, "The Subtleties of Settling Time," *The New Lightning Empiricist*, Teledyne Philbrick, June 1971.

AUTHOR'S BIOGRAPHY

Jim Williams is a staff scientist at Linear Technology Corp, where he specializes in analog-circuit and instrumentation design. He has served in similar capacities at National Semiconductor, Arthur D Little, and the Instrumentation Laboratory at the Massachusetts Institute of Technology (Cambridge, MA). A former student at Wayne State University (Detroit), Williams enjoys sports cars, art, collecting antique scientific instruments, and restoring old Tektronix oscilloscopes.

Introducing EDN's New Microsite on LEDs

Your complete source for up-to-the-minute information on LEDs and their applications.

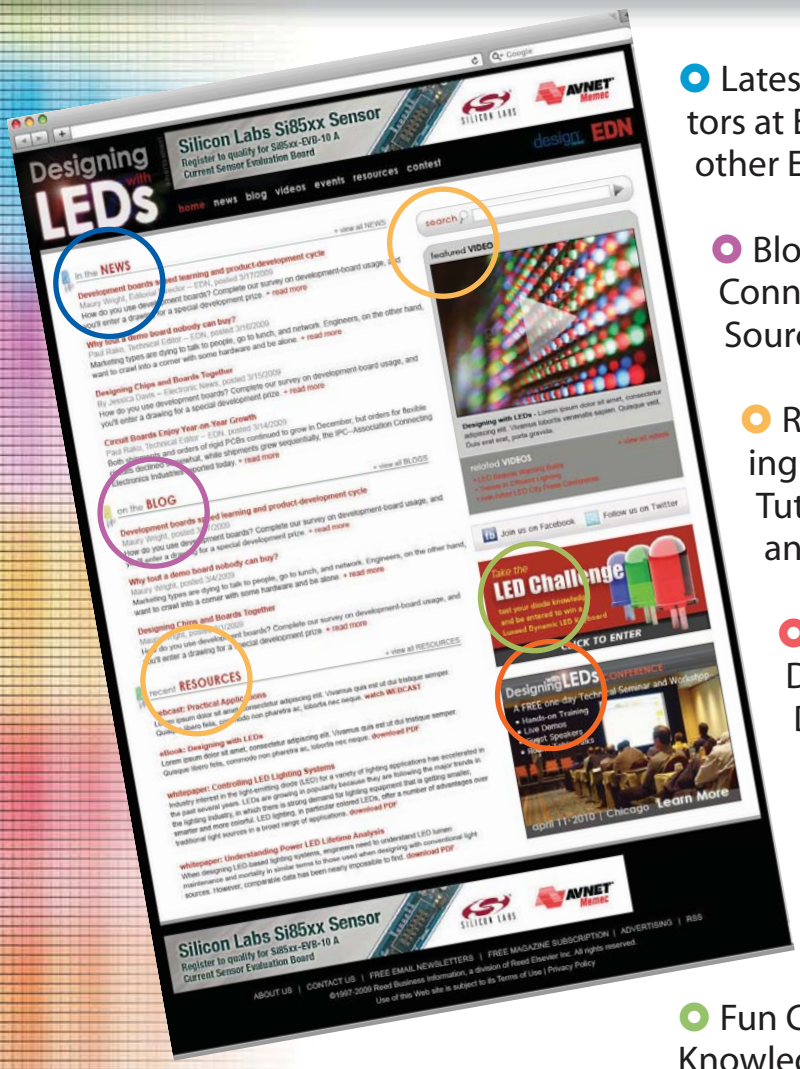

- Latest News from the Editors at EDN, Design News and other Engineering Magazines
- Blogs, including Margery Conner's Popular "Power-Source" Blog
- Resource Center, Including White Papers, eBooks, Tutorials, Webcasts and More
- Videos, Including Design Tips, Application Demos, and More
- Updates on Future In-Person Workshops and Seminars, like this one!
- Fun Challenges to Test Your Knowledge of LEDs

Check It Out Today at www.edn.com/leds

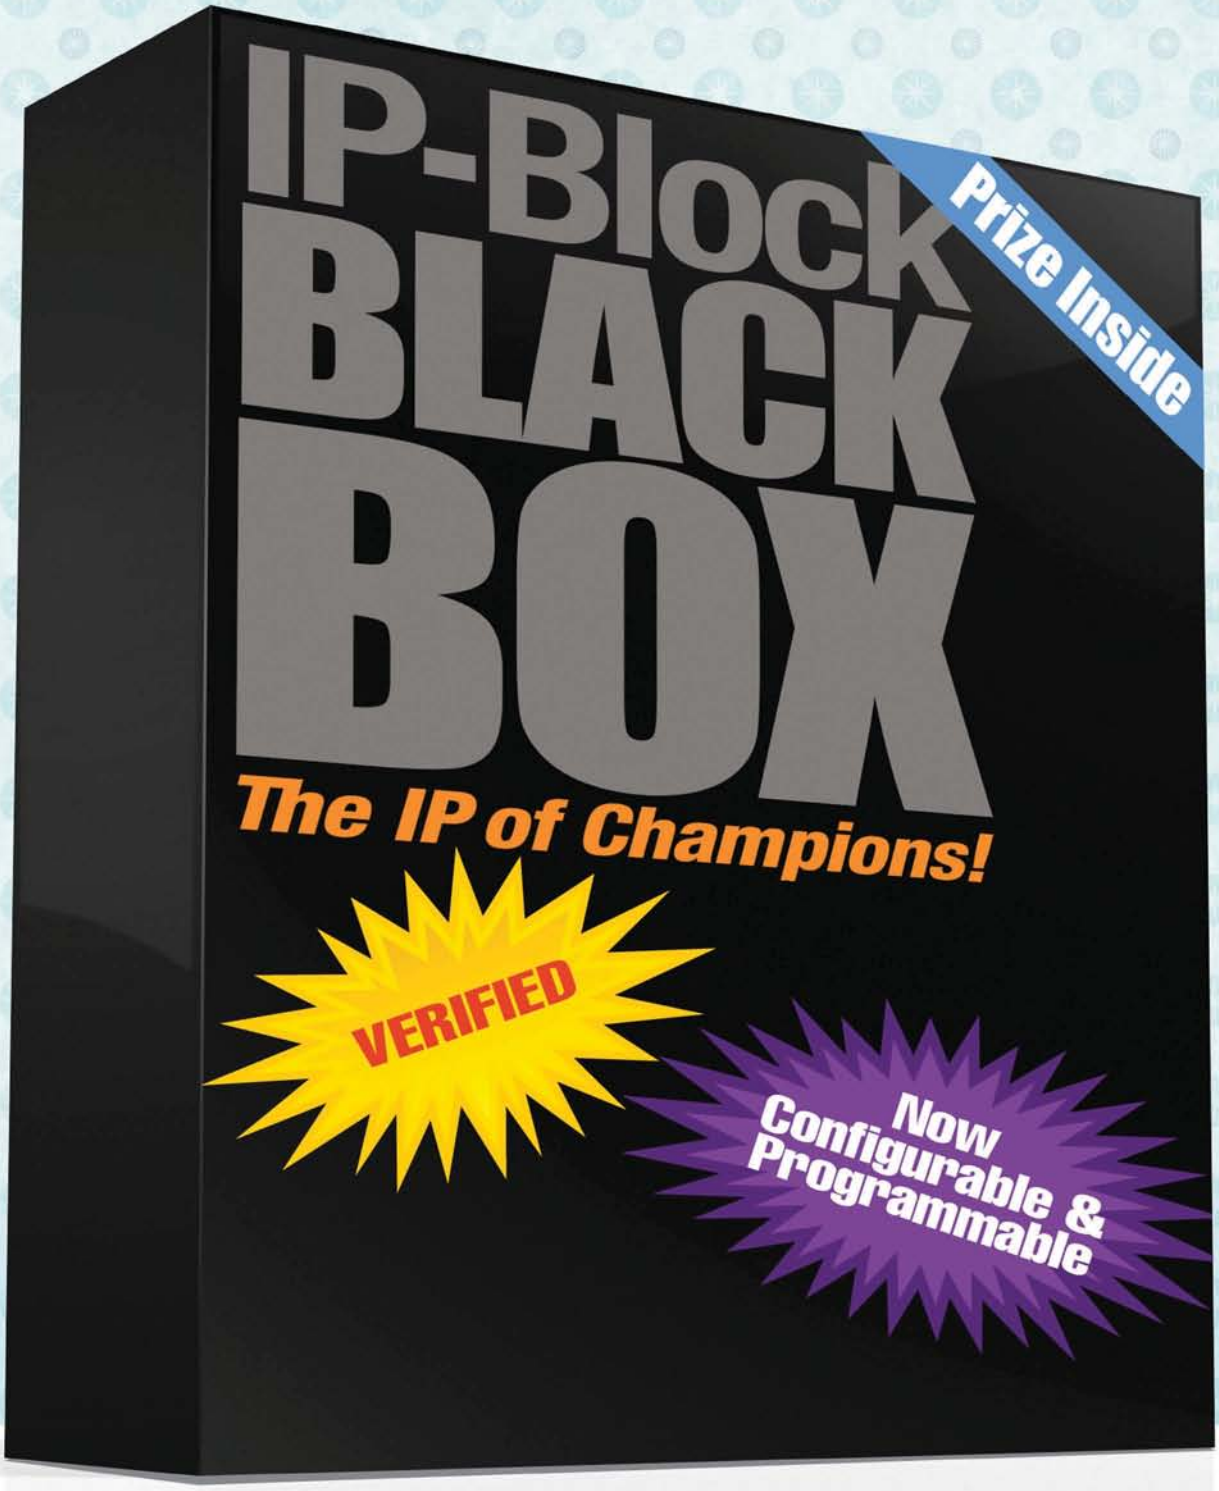
IP-Block
BLACK
BOX
Prize Inside
The IP of Champions!
VERIFIED
Now
Configurable &
Programmable

IP INTEGRATION: IS IT THE REAL SYSTEM-LEVEL DESIGN?

BY RON WILSON • EXECUTIVE EDITOR

The search for productivity in SOC (system-on-chip) design is a search for balance between abstraction and automation. Greater abstraction at a step in the design flow means fewer design elements to process. Greater automation means that each element requires less human attention. Ideally, designers could capture an abstract representation of an SOC's intended behavior, verify that the representation describes the desired chip, and push a button to tape-out. We are not yet there.

For years, some enthusiasts have promoted high-level design languages—often dialects of C—as the path to greater abstraction. Except in a few categories of architectural elements, however, it has been difficult to move the design beyond behavioral or transaction-level representation and into the implementation flow. "High-level synthesis is still very domain-specific," says Ken Wagner, PMC-Sierra's vice president of engineering. Without synthesis, designers must recode the high-level version by hand into RTL (register-transfer-level) logic.

While EDA vendors were struggling with synthesis, another kind of high-level abstraction had become common in SOC design: IP (intellectual-property) reuse. You would not ordinarily think of reusing IP as a high-level design method; it's a simple matter of expediency. The technique is so ubiquitous that many of today's SOCs simply couldn't exist without it (**figures 1** and **2**). "In our shorter designs, often 80% of the chip is composed of IP," says Jitu Kahne, director of central engineering at AppliedMicro.

Yet this heavy reliance on IP is an attempt to raise the level of abstraction in the design. Ideally, an IP block can be a black box that the design team need never open. IP reuse allows a design team to represent a large block by its function and its I/Os rather than by its internal structure. The design team could carry this abstraction all the way through to sign-off, integrating the black boxes together rather than opening them and exploring their contents.

With IP, as with high-level languages, however, life is not ideal. "IP is often a mirage," says Wagner. "Using third-party IP takes a large effort: 30 to 50% of the work required to design the block internally."

Much of this effort is due to the loss of automation, and some is due to the inability to maintain the abstraction through critical phases of the design cycle. According to Kahne, 60 to 70% of AppliedMicro's design cycles are IP integrations. Few tools, however, are available for a design

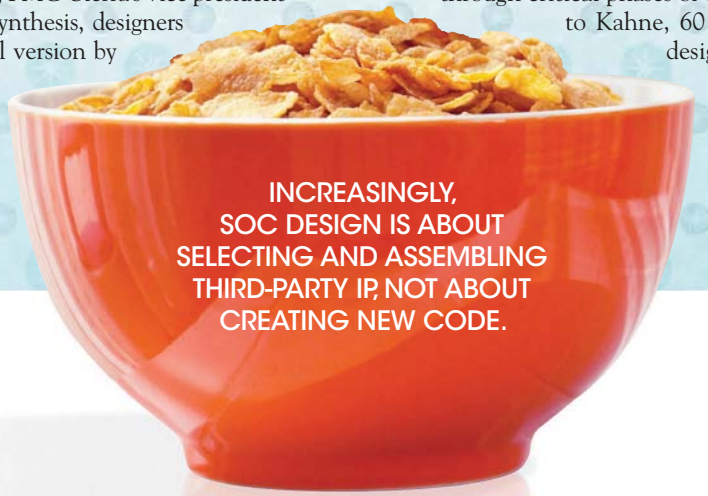

INCREASINGLY, SOC DESIGN IS ABOUT SELECTING AND ASSEMBLING THIRD-PARTY IP, NOT ABOUT CREATING NEW CODE.

flow in which selecting and integrating IP dominate.

According to Karim Arabi, senior director of engineering at Qualcomm, 80 to 90% of the IP in the company's designs doesn't change from one chip to the next. "But we still have to redo the whole integration process every time," he says. "There are no tools to automate IP qualification or integration, and the ECO [engineering-change-order] process is purely manual."

It is fair to say that IP integration dominates the SOC-design flow. Although IP reuse can greatly increase the level of abstraction, it falls short on automation. This article examines four critical activities in IP integration: determining IP behavior, constructing interfaces, filling in the infrastructure, and closing the design, comparing how designers now perform the design tasks and speculating on the opportunities for further automation.

AT A GLANCE

- IP (intellectual-property) reuse may be the best system-level chip-design method.
- Successful reuse depends on treating IP blocks as black boxes, but this approach is now rare.
- An emerging IP-integration flow overlays the traditional RTL (register-transfer-level)-to-GDSII (Graphic Design System 2) flow.
- Many opportunities exist for automation to make this flow into a methodology.
- Abstraction relies on relationships, not tools.

IP BEHAVIOR

How do designers determine the behavior of an IP block, and how do they explore the behavior of their SOCs with the blocks in place? Both questions are nontrivial if you are attempting not to look inside the black box.

A simple answer to both questions is sufficient when the IP implements a standard function, such as a USB (Universal Serial Bus) 2.0 interface. "The IEEE has done a wonderful service by standardizing so many interfaces—both external and within the die," observes Taher Madraswala, vice president of engineering at Open-Silicon. If a block implements an IEEE standard or a de facto industry standard, little of the block's behavior is subject to interpretation. And third-party VIP (verification IP) exists to confirm the block's behavior.

Not all IP implements standard functions, however. "The customer must prepare to make a significant investment to select IP to meet his specifications," PMC-Sierra's Wagner says. Some IP packages include executable behavioral or transaction-level models, which

USING AND ADOPTING IP-XACT

By Scott Wolf, Xilinx

Xilinx is adopting Accellera's IP (intellectual property)-Xact to better serve its customers in their use and reuse of IP, to provide consistency for its IP partners, and to improve internal efficiencies. IP-Xact is an industry standard, and many independent IP companies are adopting it. As such, Xilinx's adoption of the standard will enable the company's IP ecosystem to easily provide IP that is compatible with Xilinx's development tools, translating into more IP options for its customers. The company's conversion to IP-Xact includes IP and software tools, which Xilinx will implement in phases.

The first phase will support logic-based IP and tools centering on the Xilinx Core generator system, which provides Xilinx customers with a catalog of architecture-, domain-, and market-specific IP as part of the ISE (integrated-synthesis-environment) design suite. Xilinx is using the IP to capture IP-specific information, including file and file grouping semantics. This information includes synthesis files, such as sample design files and IP documentation; interface information, such as ports, generics, and their types; hardware compatibility; catalog location; and the like. Xilinx is also using the standard to control the generation and capture of configuration settings of instantiated IP cores. The conversion to IP-Xact is twofold: The company needs to convert not only Xilinx IP, but also the software tools to take advantage of this information. Xilinx has allocated a substantial amount of resources to creating supporting tools, libraries of software, and processes employing IP-Xact.

As the company implements the standard in its IP and tools, it has encountered some challenges. For instance, it had to make decisions regarding the interpretation of the IP-Xact specification. These decisions include cases in which the design environment is specifically stated to solve a requirement and cases in which IP-Xact has more than one method with which to model information. The standard provides customized generators, but the availability of these generators has created a significant amount of debate over their proper implementation in a design environment.

The use of vendor extensions is another challenge. Certain programmable features that Xilinx's IP and software use to take advantage of its FPGAs must use Xilinx-specific vendor extensions in optional ports, vectored interfaces, and hardware-architecture compatibility. Although the IP-Xact file describing the IP is a convenient location in which to centralize data associated with the IP, Xilinx must make judicious decisions to determine when it is necessary to take advantage of vendor extensions to centralize this data. The downside of using vendor extensions or of making interpretation decisions about the specification is the risk of introducing incompatibilities with other IP or software tools, and that risk works against the company's goal of interoperability.

The ratification of IP-Xact as IEEE-1685 is a significant step forward. As the standard continues to evolve, Xilinx looks forward to its increasing adoption in the industry.

Author's biography

Scott Wolf is an engineering manager at Xilinx.

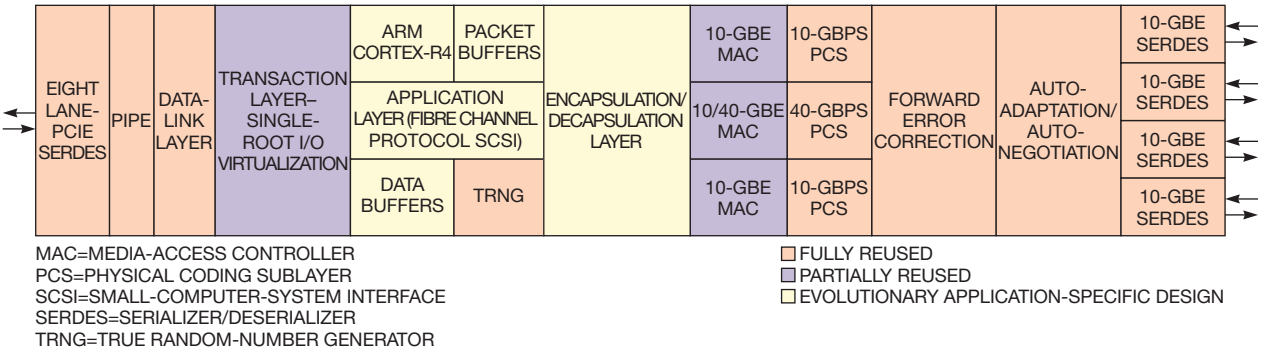

Figure 1 In this networking SOC, LSI wrapped a new packet-processing core with interface IP.

may not accurately model the behavior of the RTL logic. Without trusted models, the architects must rely on RTL simulations or FPGA emulations and on close study of data sheets to determine precisely what the IP does. To understand the workings of the full SOC with the IP in place, the architects may have to rely on mixed-mode simulation if it is fast enough, or they may have to write their own TLMs (transaction-level models) using these sources—a labor-intensive and error-inviting undertaking. This problem brings up the opportunity for automation: a tool that would extract TLMs from RTL code.

Some important considerations in modeling the behavior of IP are that most IP blocks are configurable and some, such as CPU cores, are programmable. Configurable or programmable blocks may require a good deal of configuration or even a software project to produce a high-level model of the IP for system exploration.

INTERFACING TO THE BLOCKS

Extracting the interface definition, converting it into design requirements, and verifying that the block connects correctly to the blocks with which it interacts are all significant design tasks in cases lacking an industry standard to unambiguously define the pins on an IP block. These tasks are now almost entirely manual.

As with understanding the behavior, the best case for understanding IP interfaces occurs when those interfaces comply with external standards. If the interface passes inspection by independent VIP, you should be able to just drop

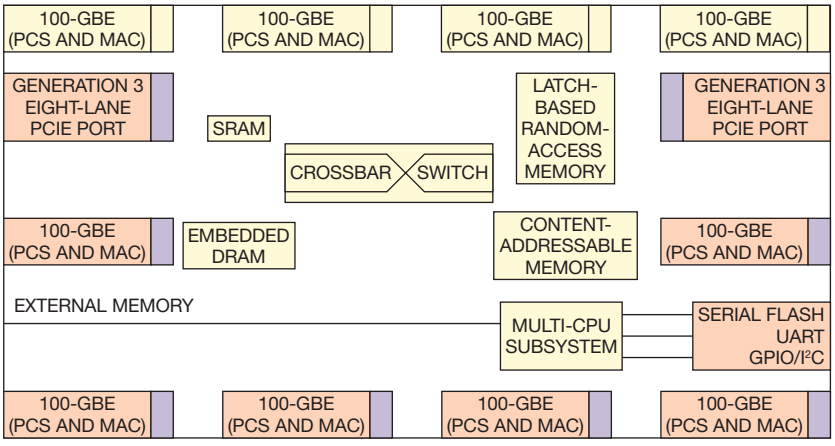

Figure 2 When design requirements change dramatically, you can reuse only the I/O ports themselves.

the black box into your design. Some FPGA tools do just that task: They build systems by assembling IP blocks onto a standard bus and by building a register map and driver library for use by software developers.

The Accellera standard IP-Xact takes a similar but more ambitious approach in the SOC world, attempting to code enough information about the function and interface behavior of an IP block into a standardized XML (Extensible Markup Language) form so that automated tools can identify an IP candidate based on design requirements and generate the necessary interface code to build the block into an SOC (see **sidebar** "Using and adopting IP-Xact"). A team in the research group at Philips Semiconductors, later NXP, developed automated system-building tools to work with these XML definitions, but the group appeared to have largely abandoned the work when NXP slashed its research activities and then divested its SOC-design teams.

The principle is sound, but, as always, there are worrisome details. "Unfortunately, not everyone follows naming conventions, even on standard interfaces," says Open-Silicon's Madraswala. So connecting the pins may require some further research and design, especially with regard to that ancient saboteur of designs: polarity. Also, even with standard interfaces, you can't blindly trust the test bench.

"It may be difficult to fully exercise the timing of interface IP, even with in-circuit emulation, unless you can create and apply appropriate stimuli and monitors to represent your use of the block," says PMC-Sierra's Wagner. His point applies even if you are staying strictly within a standard. If you are doing anything beyond the standard, such as overclocking or sideband signaling, building a custom test bench is essential. Wagner goes on to counsel engineers to ensure that the bench offers both high coverage and high performance. "Ideally, there would be static and dynamic assertions," Wagner adds.

The question of assertions is a sensitive one. More and more SOC teams are adopting assertion-based verification strategies. But an informal survey suggests that assertions are appearing in IP deliverables at a much slower rate. This absence may force the SOC team to create their own assertions for IP, which would require them to intensively mine the data sheet and probably examine the RTL, as well. There is hope from the EDA world in this regard, though. Start-ups Zocalo Tech and NextOp Software recently introduced tools that could help with creating assertions. The Zocalo tool identifies signals for which, according to its heuristics, assertions would be appropriate. The NextOp tool goes a step further by synthesizing assertions. How well either tool would work with black-box third-party IP remains to be investigated.

In a less-desirable case, a port on the IP block does not follow any well-defined standard. In this case, as in creating your own assertions, the task comes down to the designer and the data sheets. "You must analyze the semantics of the interface as it is specified in the IP data sheet," Wagner says. For complex ports this can be an error-prone and grindingly manual process.

Understanding how the ports on an IP block work is part of the challenge. Connecting them into the SOC design is a separate task that may require additional logic design—FIFO (first-in/first-out) buffers or protocol translators, for example. No synthesis tool can generate these interfaces. Ideally, as the IP block merges into the SOC, its test bench can contribute to the overall SOC-verification suite, generating stimuli and assertions for transactions that cross the boundaries of the block. But this reuse of VIP is also a manual task begging for automation.

THE DETAILS DEVIL

Beyond the basic function of an IP block and the interfaces that connect it into the SOC, there are still more integration issues—for example, power management. SOC architects now have a bulging arsenal of techniques for reducing dynamic and static power at the chip level. These approaches include device-level techniques, such as a choice of transistors with different threshold voltages or with adjustable back biases to change their thresholds; circuit-level techniques, such as signal gating and fine-grained clock gating; and block-level tools, such as coarse-grained clock gating, voltage islands, DVFS (dynamic voltage-frequency scaling), and power gating. The issue for IP integration is how to coordinate

FOR MORE INFORMATION

Accellera
www.accellera.org

AppliedMicro
www.appliedmicro.com

LSI Corp
www.lsi.com

NextOp Software
www.nextopsoftware.com

NXP
www.nxp.com

Open-Silicon
www.open-silicon.com

PMC-Sierra
www.pmc-sierra.com

Qualcomm
www.qualcomm.com

Xilinx
www.xilinx.com

Zocalo Tech
www.zocalo-tech.com

the power-management strategy at the chip level with whatever assumptions the IP designers have made.

"It helps if the chip architects understand the power-management provisions in the IP blocks," Madraswala says. "Can the block shut itself down? Does it need to be an independent power island? Is it capable of DVFS?" Designers should fold all this data-sheet information about the capabilities of the IP blocks into the power planning for the whole chip.

Once the project moves from planning to implementation, the influence may flow in the other direction. "For soft IP, the customer generally sets the strategy for power," Wagner says. "Implementation work may fall to the IP vendor or the customer." The point of using IP in the first place is to achieve a higher level of abstraction for the design team. So opening up the RTL to weave in a new power-management strategy is not a happy alternative. However, the chip-design team may sometimes be able to improve the power consumption of a block without understanding its structure.

Synthesis tools implement netlist-level power-saving techniques, such as signal or clock gating, requiring little work by the chip-design team. The team should always use logic-equivalence checkers, cautions Wagner, when modifying netlists with which they are unfamiliar. It's easy for a tool to alter performance along with efficiency.

THE POINT OF USING IP IS TO ACHIEVE A HIGHER LEVEL OF ABSTRACTION FOR THE DESIGN TEAM.

Approaches that alter the structure of the IP, such as power gating, may require either the active participation of the IP developer or reverse-engineering by the SOC team. Such changes may also require alterations to the test bench just to test the power-saving modes. Design-management-team members may have to balance how much energy they can save in a block against just how much abstraction they are willing to surrender to get the savings.

Scan insertion, another downstream automated task, should create no obstacles to black-box treatment of IP. Exceptions may exist, however. "High-performance-IP vendors tend to be very conservative about scan chains in their blocks," Madraswala says. "They don't want scan flip-flops on their critical paths." Madraswala suggests checking the fault coverage you are getting rather than blindly trusting the vendor's scan directives.

Madraswala also says that clock insertion should not require the chip designers to get too involved in the entrails of an IP block. All the necessary information to generate the clock structures should be in the deliverables, and the clock requirements should be

in the data sheet. In this area at least, the box should remain black, as long as you don't reorganize the clocks to accomplish coarse-grained clock gating.

DESIGN CLOSURE

Ideally, you could close the design—confirming timing, signal integrity, power integrity, and design-rule compliance—without looking into the IP blocks. The IP vendor presumably has stated that the block will pass static-timing analysis at the specified frequency, will be free of signal-integrity issues, and will be DRC (design-rule-checking)-clean. Because most IP today uses fully registered inputs and outputs, there should be no timing paths running across the boundaries of the block. So it should be no problem to hierarchically close timing.

The IP vendor bases its confidence on its own synthesis, placement, and routing using its own tools, scripts, and libraries, however. As the proverbial fine print says, your results may vary. "It may be necessary to resynthesize the IP to ensure it will meet timing using the customer's standard-cell and SRAM libraries," Wagner says.

The interesting question is what happens if something fails. If the design team has been successful in treating the IP block as a black box, they will be on unfamiliar ground in trying to resolve closure issues. The only way to preserve the integrity of the box will be to involve the IP developer in resolving closure issues. The developer, presumably, understands the netlist and the constraint files. If a novel use of the IP or changes to support new power-management modes also make the situation unfamiliar to the developer, then there may be no alternative to collaboration between the IP- and SOC-design teams.

Even within an organization, advanced power management can force close cooperation between IP and chip teams. "We use power gating in some situations," says Harmel Sangha, senior director of marketing at LSI. "When we do, we depend on interpersonal relationships and regular reviews to convey our intent."

"It really helps that these designs involve basically the same guys each time," says Robert Madge, the company's director of technical marketing. "If we are working with third-party IP, we will bring the supplier into our team. For example, we may put their IP on our own test chip and analyze the results with them."

Surveying some critical steps in the IP-integration process, you can see that, in some cases, new tools are helping with automation. In other cases, there are opportunities for new tools to greatly ease the burden of the integration flow. The only way to preserve the abstraction of IP as a black box, however, may be to build an intimate, engineer-to-engineer collaboration between the IP-creating team and the SOC team. No tools are available to automate that relationship. EDN

BY PAUL RUTHERFORD • ISILON SYSTEMS

Data storage yields increased design productivity

REVIEW HOW CURRENT DATA-STORAGE TECHNOLOGIES CAN OFFER AN ESSENTIAL TOOL FOR SIGNIFICANTLY INCREASING THE SPEED OF DESIGN PROCESSES, ENHANCE PRODUCTIVITY, AND AVOID THE UNPREDICTABILITY OF EVOLVING DESIGN NEEDS.

As IC designers continue to look for ways to improve design productivity, one area that deserves scrutiny at 90 nm and below is the storage environment for design files. Just as the CPU, networking, and EDA tools all play a role in boosting productivity, so does the storage environment. Current data-storage technologies can offer design teams essential tools for significantly increasing the speed of design processes. The storage that supports design workflows is significantly more complex than it may at first appear, and, as designers move toward 65-nm and smaller process technologies, those complexities increase. As design leaders and their partners in IT evaluate and integrate storage environments, they must consider key criteria to account for the efficacy and budget impact of storage in the design flow.

These criteria include I/O performance—that is, the speed of writing data to disk and reading it back to the CPU and user. Designers can further break down this criterion into transactional, sequential, and random-access patterns.

Other criteria to consider include availability, the cost of maintaining uninterrupted access to the data; reliability, the cost of mitigating the risk of data loss; and maximum scalability, the largest data volume the storage system can ever hold. You might also want to consider dynamic configuration, the cost in time and effort it takes to make a change to the system, and resolution of scale, the smallest increment by which you can increase the storage capacity. Also important are the purchase price; the total cost of ownership—that is, how much it costs to operate the system over its useful lifetime; and the system's ease of use.

As the central point of a one-to-many operation in the design flow, storage acts as the intermediary between the application and the results, back annotations, libraries, IP (intellectual property), team members' data "sandboxes," test suites, and design repositories. Storage has emerged as a barrier to productivity at processes below 90 nm. The reason for this slowdown in productivity is a dramatic increase in the number of files and rule decks necessary for verification and in the number of libraries and cells in designs. In addition, the overall gate counts continue to grow.

These changes have created two challenges that consistently affect designers. First, designers require fast file I/Os to bring files and rule decks to the EDA tool; otherwise, the tool remains idle. Second, the storage environment must scale in its available capacity; otherwise, factors such as intermediate formats, "scratch" space, and the overall size of design databases will overrun the storage space available to the tools. In short, both performance and capacity are essential for avoiding storage-related bottlenecks. Depending on the storage architecture you use, you may find it difficult to simultaneously achieve both of these goals.

STORAGE ACTS AS THE INTERMEDIARY BETWEEN THE APPLICATION AND THE RESULTS, BACK ANNOTATIONS, LIBRARIES, IP, AND MORE.

File-I/O performance and scalability have taken a lot of designers by surprise because they previously encountered storage environments that were adequate for the task. For example, traditional storage employing RAIDs (redundant arrays of independent disks) has a capacity limit of 16 Tbytes in a file system (**Figure 1**). Five years ago, 16 Tbytes represented a lot of space; today, it is a fraction of what a designer might need. It is common to see a 45-nm design with a total design database of more than 200 million files, totaling 30 Tbytes, whereas, at 32 nm, the design would require 100 Tbytes.

For designs that now consume a few terabytes, designers should consider mapping a storage strategy that accounts for future growth. Scripting, run directories, shared design repositories, and intermediate-results files all can become time-consuming and error-prone without a storage architecture

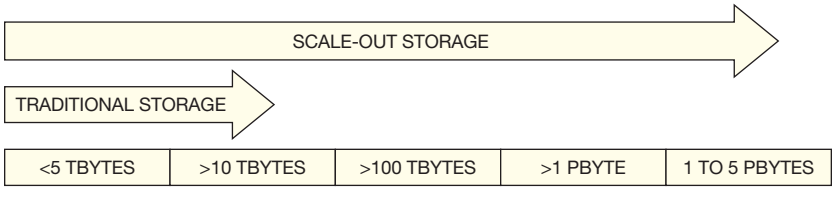

Figure 1 Traditional storage employing RAIDs has a capacity limit of 16 Tbytes in a file system.

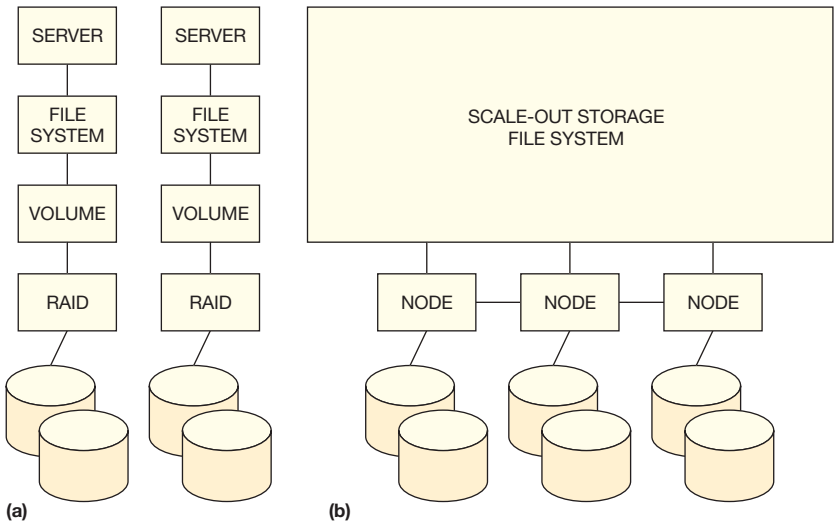

Figure 2 With traditional RAID-based storage systems, a filer head unit manages the traffic into and out of the array of disks (a). This architecture involves a single file system on server maps to RAID. It works only in CPUs in server head units and allows no sharing across servers. In a one-file-system approach, capacity is approximately 5 Pbytes (b). It abstracts multiple storage devices, integrating data protection and storage management. Its performance expands when you add nodes, and all nodes can share these performance enhancements.

that allows the file system to scale to accommodate these increasingly massive file repositories.

File-I/O performance is a separate concern and requires even more effort to manage. With traditional RAID-based storage systems, a filer "head" unit manages the traffic into and out of the array of disks (**Figure 2**). This approach makes it easy to add additional disks to expand capacity, but a "hot-filer" problem ultimately arises when the filer head reaches its performance limit. The impact of this scenario is most noticeable in physical verification. Here, the number of rules and libraries that the design must access at processes of 65 nm and below can overwhelm a traditional storage system. In one case, the ratio of elapsed time to CPU time during DRC (design-rule checking) was roughly 20-to-1; a hot filer had made storage-I/O performance the bottleneck in the physical-verification runtime. During both physical implementation and physical verification, a good way to confirm whether file I/O is a bottleneck is to use the EDA-tool-log files to capture this information. At ratios of elapsed time to CPU time of greater than 3-to-1, file-I/O issues are most likely the gating factors for achieving results.

PLANNING STORAGE FOR PRODUCTIVITY

Storage-capacity planning for IC design is difficult in that requirements change rapidly and irregularly. Planning for growth according to the gate count or number of designers is often insufficient when, for instance, the use of multiple-threshold-voltage libraries can double the amount of scratch space an EDA tool requires. Similarly, a revolutionary new physical-verification technique might increase needed I/O performance by an exponential amount. To be responsive to the requirements of IC designers, an ideal storage architecture should be scalable in performance and capacity in both small and large increments without requiring a system redesign or replacement. Ideally, a storage approach should have pay-as-you-grow characteristics that allow the organization to acquire additional resources only as need dictates, rather than overproviding storage space as a means of future-proofing.

Some industries deal with large data sets and high I/O performance from which IC designers can learn as their designs reach similar levels of complexity. For example, in the life sciences, DNA (deoxyribonucleic-acid)-sequencing processes typically create results files that are a few terabytes each and that require iterative analysis to interpret correctly. Media production, especially for films and television, also generates files in terabytes that require high I/O performance to edit and render in a reasonable amount of time. From studying how these industries have dealt with their performance and scalability issues, designers will find some practices that directly apply to IC design.

First, the ability to quickly add capacity without restructuring the data across multiple file systems saves significant downtime as design requirements change. In these cases, scale-out storage technologies, in which you can add capacity without downtime or reconfiguration, are attractive. Traditional RAID-based storage systems require downtime to add capacity. Even if the downtime takes place during off-peak hours, it prevents the ability to perform critical functions, such as overnight runs. If the added storage capacity exceeds 16 Tbytes, you must also absorb the time and complexity of restructuring which data belongs in which file system, and you must adapt the tool scripts to reflect that restructuring.

You must also ensure that you can scale capacity and performance independently of each other. In some cases, a design team simply needs more I/O performance but doesn't want the expense of adding capacity, and vice versa. Traditional RAID-based storage limits your ability in this regard because of the architecture of the filer head to the disks. An alternative to that approach is to use an architecture such as scale-out storage, which distributes the processors and memory that define I/O performance across all nodes of the storage environment instead of just the filer head unit. This approach allows performance increases by adding nodes that contain strictly processor and memory but no hard drives. This approach works well if your goal is to add just performance without adding capacity. In addition to giving the design team exactly the I/O performance it needs to remove the storage bottleneck to tool performance, this approach removes the cost overhead of additional disks.

A third practice is to make ease of use an explicit goal of using a new storage environment. Ease of use applies to man-

agement overhead, the effort necessary to scale, and the impact on the user. For management, the human resources necessary for maintaining a large storage system range from none to many full-time employees. The management of an appropriately sized, well-built storage system should not require the hiring of additional dedicated IT staff.

Scaling a storage system's capacity, performance, or both, whether by fractional amounts or by orders of magnitude, should not require multiple months of meetings to plan or even several days of IT expertise to implement. Designers should instead be able to scale a storage environment in minutes, independently of scale. Ideally, IC designers focus on the design, not computers or disks. Designers shouldn't be concerned with or aware of volumes, capacity, formats, or how to access their data. A user might notice a sudden increase in capacity but never experience an interruption in service.

Storage decisions for IC designers typically involve four primary questions: Will the storage approach provide sufficient performance and capacity for the designers' current needs? Will the system experience any significant downtime or data loss? Does the company have the human resources to manage the system? How long will it be before the company needs to upgrade this system and at what cost? Today's traditional file-based storage approach cannot cost-effectively manage the rapid data growth taking place in IC design and, as such, can become an impediment to efficient IC design.

TODAY'S TRADITIONAL FILE-BASED STORAGE APPROACH CANNOT COST-EFFECTIVELY MANAGE THE GROWTH TAKING PLACE IN IC DESIGN AND CAN BECOME AN IMPEDIMENT TO EFFICIENT DESIGN.

By deploying scale-out technologies, specifically scale-out storage for file-based data, IC design teams can take advantage of the pay-as-you-grow scalability of both performance and capacity, a single-file system, and the inherent ease of use to meet the demands of IC design today and in the future. By understanding the differences in storage architectures and by planning for future design needs, IC-design managers and their IT teams can craft a storage strategy that avoids design-productivity pitfalls and minimizes IT-related delays in expanding both storage performance and capacity, accelerating the IC-design process and improving time to market. EDN

AUTHOR'S BIOGRAPHY

Paul Rutherford is chief technology officer at Isilon Systems. He has more than 15 years of experience in guiding long-term strategic-technology visions. Before joining Isilon, Rutherford was chief technology officer at ADIC and served at EMass, where he led software development for high-capacity storage systems. He received a bachelor's degree in computer science from the University of Minnesota (Minneapolis).

16-Bit, 1Msps SAR ADC with 94dB SNR

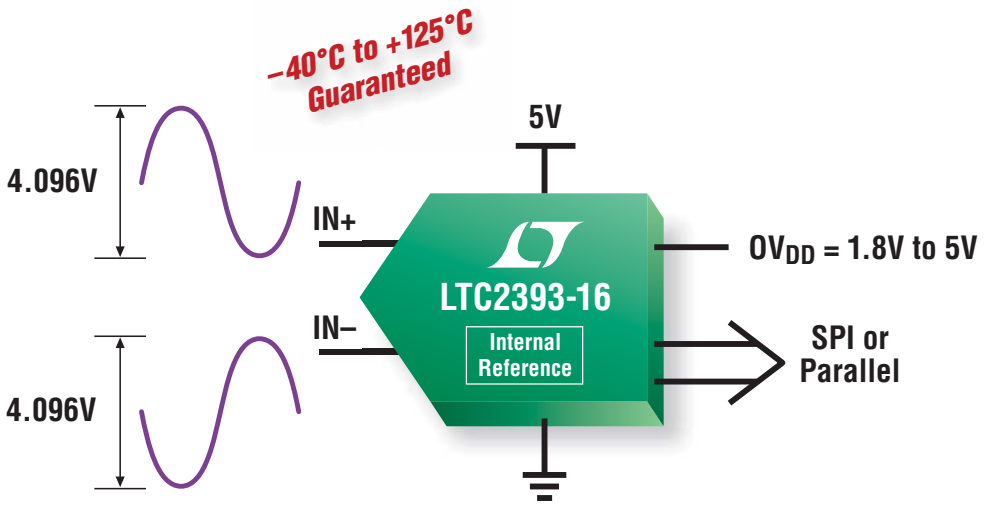

±4V$_{P-P}$ Differential Input Span ADC Excels in Noisy Environments

The LTC®2393-16's wide differential input range of ±4.096V and outstanding 94dB dynamic range ensures sufficient noise margin in hostile industrial environments. Operating from a single 5V analog supply, this ADC achieves 16-bit no missing code resolution at 1Msps with no latency. The LTC2393-16 includes an internal 20ppm/°C reference and is ideal for medical instruments, spectrum analysis, ATE, instrumentation and high-resolution data acquisition.

Features

- Pin- and Software-Compatible Family:
  - LTC2393-16: 1Msps
  - LTC2392-16: 500ksps
  - LTC2391-16: 250ksps
- Single 5V Analog Supply, 1.8V to 5V Digital Supply
- ±2LSB INL (Max), ±1LSB DNL (Max)
- 20ppm/°C (Max) Internal Reference
- –40°C to +125°C Guaranteed Temperature Range
- Parallel & Serial Interfaces
- 48-Pin 7mm x 7mm QFN & LQFP Packages

16k Point FFT f_S = 1Msps, f_{IN} = 20kHz

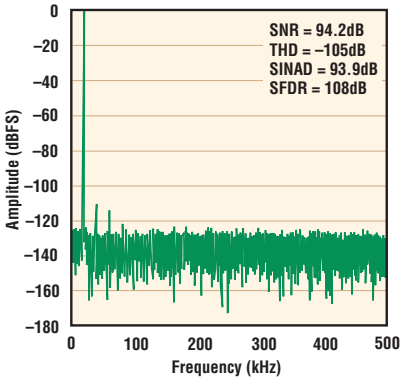

Info & Free Samples

www.linear.com/2393

1-800-4-LINEAR

Free Industrial Signal Chain Brochure

www.linear.com/indsolutions

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Simple battery-status indicator uses two LEDs

Abhijeet Deshpande, People's Education Society Institute of Technology, Bangalore, India

Properly maintained rechargeable batteries can provide good service and long life. Maintenance involves regular monitoring of battery voltage. The circuit in **Figure 1** works in most rechargeable batteries. It comprises a reference LED, LED_{REF}, which operates at a constant current of 1 mA and provides reference light of constant intensity regardless of battery voltage. It accomplishes this task by connecting resistor R_1 in series with the diode. Therefore, even if the battery voltage changes from a charged state to a discharged state, the change in current is only 10%. Thus, the intensity of LED_{REF} remains constant for a battery state from a fully charged state to a fully discharged state.

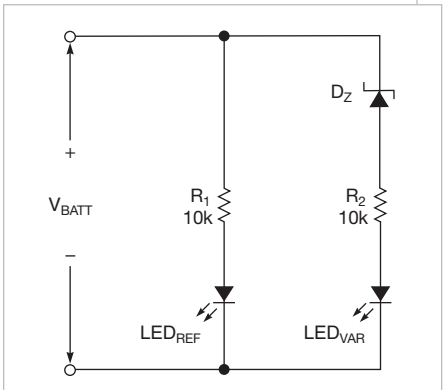

Figure 1 This circuit works in most rechargeable batteries. It comprises a reference LED, LED_{REF}, which operates at a constant current of 1 mA and provides reference light of constant intensity regardless of battery voltage.

The light output of the variable LED changes with respect to changes in battery voltage. The side-by-side-mounted LEDs let you easily compare light intensities and, thus, battery status. Using diffused LEDs as crystal-clear LEDs can damage your eyes. Instead, mount the LEDs with sufficient optical isolation so that the light from one LED does not affect the intensity of the other LEDs.

The variable LED operates from 10 mA to less than 1 mA as the battery voltage changes from fully charged to fully discharged. Zener diode D_Z in series with resistor R_2 causes the current to change with battery voltage. The sum of the zener voltage and the drop across the LED should be slightly less than the lowest battery voltage. This voltage appears across R_2. As the battery voltage varies, it produces a large variation of current in R_2. If the voltage is approximately 1V, then 10 mA will flow through LED_{VAR}, which is much brighter than LED_{REF}. If the voltage is less than 0.1V, then the light intensity of LED_{VAR} will be less than LED_{REF}, indicating that the battery has discharged.

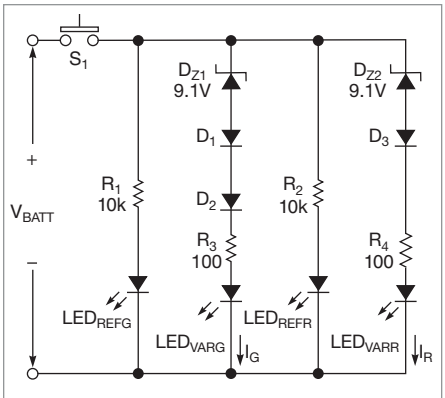

Figure 2 This circuit can withstand 13V because it has a 10-mA margin. If the LEDs are bright, quickly release pushbutton switch S_1.

Immediately after the battery has charged, the battery voltage is more than 13V. The circuit can withstand this voltage because it has a 10-mA margin. If the LEDs are bright, quickly release pushbutton switch S_1 to avoid damage to the LEDs (**Figure 2**).

The **figure** uses a 12V lead-acid battery indicator as an example, but you

TABLE 1 LED INTENSITY

Light output of LED_{VARG}	Light output of LED_{VARR}	Battery status (%)
Much brighter than LED_{REFG}	Much brighter than LED_{REFR}	70 to 100
Equally as bright as LED_{REFG}	Much brighter than LED_{REFR}	60
Off	Brighter than LED_{REFR}	50 to 30
Off	Equally as bright as LED_{REFR}	20
Off	Off	0 to 10

DIs Inside

▶To see all of *EDN*'s Design Ideas, visit www.edn.com/designideas.

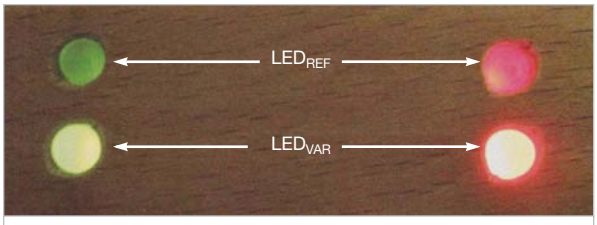

Figure 3 Both variable-intensity LEDs are brighter than the reference LEDs, indicating that the battery is 100% charged.

can extend the design to accommodate other types of chargeable batteries. You can also use it for voltage monitoring. It uses two green LEDs to indicate whether the battery has charged above 60%. A set of red LEDs indicates whether the battery charge drops below 20%. LED_{REFG} and LED_{REFR} feed through 10-kΩ resistors R_1 and R_2. For the variable-intensity LEDs, a zener diode works in series with 100Ω resistors R_3 and R_4. Diodes D_1, D_2, and D_3 provide the required clamping voltages. **Table 1** shows how LED intensity indicates battery charge.

The following **equation** calculates the variable intensity for the green LED: $V_{BATT}=I_G\times100+V_{D1}+V_{D2}+V_{LEDG}+V_{DZ1}$. For a green-LED current of 1 mA, $V_{BATT}=10^{-3}\times100+0.6+0.6+1.85+9.1=12.25V$. The selected LEDs have a drop of 1.85V at 1 mA.

If the LED has different characteristics, then you must recalculate the resistor values. At this voltage, the LEDs have the same intensity, and the battery is 60% charged. See **Reference 1** for lead-acid-battery voltages.

The following **equation** calculates the variable intensity for the red LED: $V_{BATT}=I_R\times100+V_{D3}+V_{LEDR}+V_{ZD2}$. For a green-LED current of 1 mA, $V_{BATT}=10^{-3}\times100+0.6+1.85+9.1=11.65V$.

YOU CAN USE THE DESIGN TO ACCOMMODATE OTHER TYPES OF CHARGEABLE BATTERIES. YOU CAN ALSO USE IT FOR VOLTAGE MONITORING.

At this voltage, both red LEDs have equal intensities, and the battery is 20% charged. LED_{VARG} is off. **Figure 3** shows that both variable-intensity LEDs are brighter than the reference LEDs, indicating that the battery is 100% charged. EDN

REFERENCE

1 "Lead Acid Battery Charging," *Solar Navigator*, 2005, www.solarnavigator.net/battery_charging.htm.

Hardware watchdog timer accepts range of frequencies

Robert Most, Ferris State University, Big Rapids, MI

Microcontrollers usually require a watchdog timer to bring high currents or high voltages into a safe condition. Many microcontrollers have built-in watchdog timers for this purpose. You may, however, prefer that an external circuit be the judge. An external watchdog circuit checks the integrity of the microcontroller's code, bringing the outputs to a safe state when it judges that the microcontroller's firmware has gone awry. Conditions such as errant code or failed hardware can cause these issues.

You have many options for designing an interface between your external timer and the microcontroller. For example, you can use a port pin, which yields a flexible and adaptable way to couple with almost any microcontroller.

In applications that require energizing an output or series of outputs, you need an enable signal. You can use this

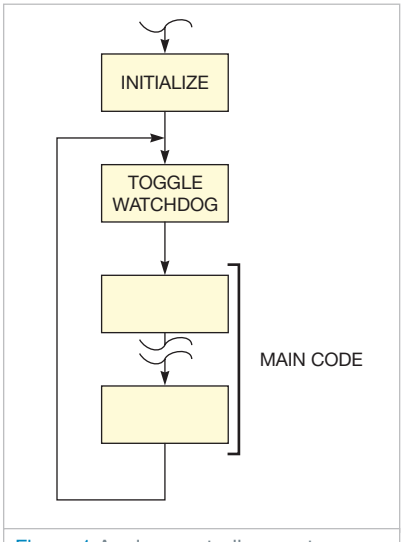

Figure 1 A microcontroller must constantly generate a pulse to keep a watchdog timer from timing out.

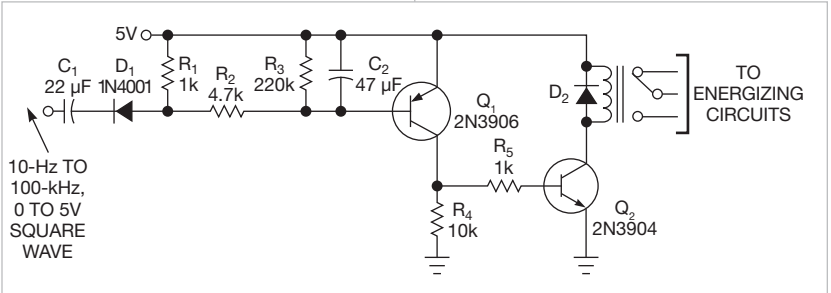

Figure 2 Capacitor C_2 remains charged enough to keep the circuit energized as long as a pulse stream of at least 10 Hz is present. If the pulse drops, the circuit removes power through the relay.

Extend battery life with the MAXQ610 16-bit microcontroller

Up to 12 MIPS in under 4mA—industry's highest MIPS/mA!

The MAXQ610 microcontroller is designed for low-cost, high-performance, battery-powered applications. This 16-bit, RISC-based microcontroller has a wide operating range (down to 1.7V) for long battery life and ultra-low power consumption. Its anticloning features and secure MMU enable you to protect your IP.

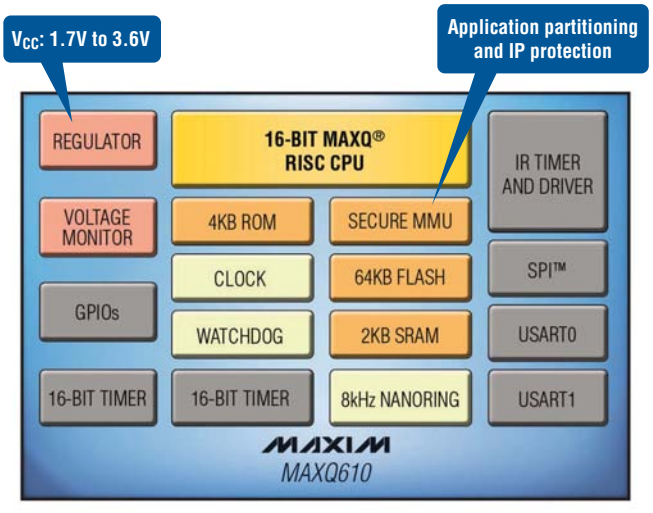

Microcontroller

- 16-bit MAXQ RISC core
- 64KB flash memory, 2KB SRAM
- Ultra-low supply current
  - Active mode: 3.75mA at 12MHz
  - Stop mode: 200nA (typ), 2.0µA (max)
- Wide, 1.7V to 3.6V operating voltage range
- IP protection
  - Secure MMU supports multiple privilege levels, helps protect code from copying and reverse engineering

Peripherals

- Two USARTs and one SPI master/slave communication port
- Two 16-bit timers/counters
- 8kHz nanoring functions as programmable wakeup timer
- IR timer simplifies support for low-speed infrared communication
- IR driver capable of 25mA sink current

Part	Temp Range (°C)	Program Memory	Data Memory	Operating Voltage (V)	Package	Price† ($)
MAXQ610B	0 to 70	64KB flash	2KB SRAM	1.7 to 3.6	40-TQFN	1.45

MAXQ is a registered trademark of Maxim Integrated Products, Inc.
SPI is a trademark of Motorola, Inc.
†1000-up recommended resale. Prices provided are for design guidance and are FOB USA. International prices will differ due to local duties, taxes, and exchange rates. Not all packages are offered in 1k increments, and some may require minimum order quantities.

www.maxim-ic.com/MAXQ610-info

MAXIM DIRECT.
www.maxim-ic.com/shop

AVNET electronics marketing
www.em.avnet.com/maxim

MAXIM INNOVATION DELIVERED®

For free samples or technical support, visit our website.

enable signal to energize a relay providing power to an output device. Turning the enable signal into a watchdog signal is therefore a wash in terms of port-pin usage. The implementation of this watchdog circuit requires the output port pin to toggle, rather than stay at a constant state.

Most microcontroller code has a main loop that is always either performing a task or calling larger functions and interacting with interrupt-service routines. If an errant task, unforeseen bug, or unintentional vector executes, the main loop either gracefully reinitializes or becomes stuck indefinitely. Either situation breaks the execution of the main loop. When that situation occurs, the timer needs to remove power from your circuits.

You can implement the watchdog timer by toggling a port output whenever the main loop runs, provided that the main loop executes 10 to 100,000 times/sec. **Figure 1** demonstrates this concept. You can implement this watchdog philosophy in several ways. You must allow for variability in the rate of the toggle signal because extraneous interrupts and other nondeterministic events may cause variability in the main loop's execution loop time. If the watchdog circuit is not sufficiently forgiving, it may lead to false triggers, defeating its purpose. The recovery time of the watchdog timer is the maximum time between toggling events in duration. This scenario can happen when the system is in a recovery, or "limp-home," mode. This circuit can accommodate the recovery duration, but, if it deems necessary, the recovery operation can disable the watchdog timer's output. You can download **Listing 1**, a document containing sample code, from the Web version of this Design Idea at www.edn/100812dia.

THE IMPLEMENTATION OF THIS WATCHDOG CIRCUIT REQUIRES THE OUTPUT PORT PIN TO TOGGLE, RATHER THAN STAY AT A CONSTANT STATE.

The circuit in **Figure 2** uses two bipolar transistors. The second transistor is simply a relay driver. The circuit works by removing the dc component from the incoming toggling square wave, rectifying and creating an average dc value. This wave feeds transistor Q_1, biased such that a prolonged absence of pulses turns it off, thus turning off Q_2 and the relay. Changing the value of C_2 also changes how quickly the watchdog timer reacts to an absence of pulses. The circuit accepts toggling frequencies of 10 Hz to 100 kHz.

Changing the value of capacitor C_2 changes the range of compliance frequencies that the main loop generates. The circuit also functions with a 3.3V processor, but either the relay must have its own 5V supply or you must use a lower-voltage relay. EDN

Get four colors from 2 bits

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

Three-color LEDs contain red, green, and blue LEDs in one package. Using two digital control signals, you can drive these LEDs to produce four colors. The circuit in **Figure 1** uses an Analog Devices (www.analog.com) ADG854 dual analog 1-to-2 demultiplexer that lets you select the current through each LED.

The circuit uses a distinct current, I or 2I, to drive each LED. The demultiplexers determine the routes of the currents through transistors Q_1, Q_2, and Q_3 in transistor array IC_2 to the LEDs. These transistors act as both current sources and summing elements.

The following **equation** yields the value of the current: $I=(V_{REF}-V_{BE})/R_E$, where V_{BE} is the base-emitter voltage of bipolar transistors Q_1, Q_2, and Q_3. The base-emitter-voltage value varies slightly depending on the total collector current, but you can neglect this variation. Refer to the data sheet of your transistor array for this information.

One unit of current constantly flows through the green LED. Demultiplexer D_1 routes another unit of current to either the red LED or the blue LED, and D_2 routes the third unit of current to either the green LED (2I total) or the red LED.

Table 1 shows the states and colors that this circuit produces. The sum of currents flowing through all LEDs is 3I at one time for all four combinations of control variables. Thus, the generated light is approximately of the same intensity regardless of color.

The decreasing value of the base-emitter voltage with temperature, which is approximately –1.42 mV/°C, causes an increase in current through the LEDs by approximately 0.33%/°C. It has a ben-

THE DEMULTIPLEXERS DETERMINE THE ROUTES OF THE CURRENTS THROUGH TRANSISTORS Q_1, Q_2, AND Q_3 IN TRANSISTOR ARRAY IC_2 TO THE LEDS.

TABLE 1 DISTRIBUTION OF CURRENT AND COLORS

IN_1	IN_2	I_R	I_G	I_B	Color
0	0	I	I	I	White
0	1	Off	2I	I	Aqua
1	0	2I	I	Off	Red-orange
1	1	I	2I	Off	Yellow

Turn it up!

Maxim's filterless Class D amplifiers provide the output power you need, and our leading emission-suppression technologies mean no EMI hassles.

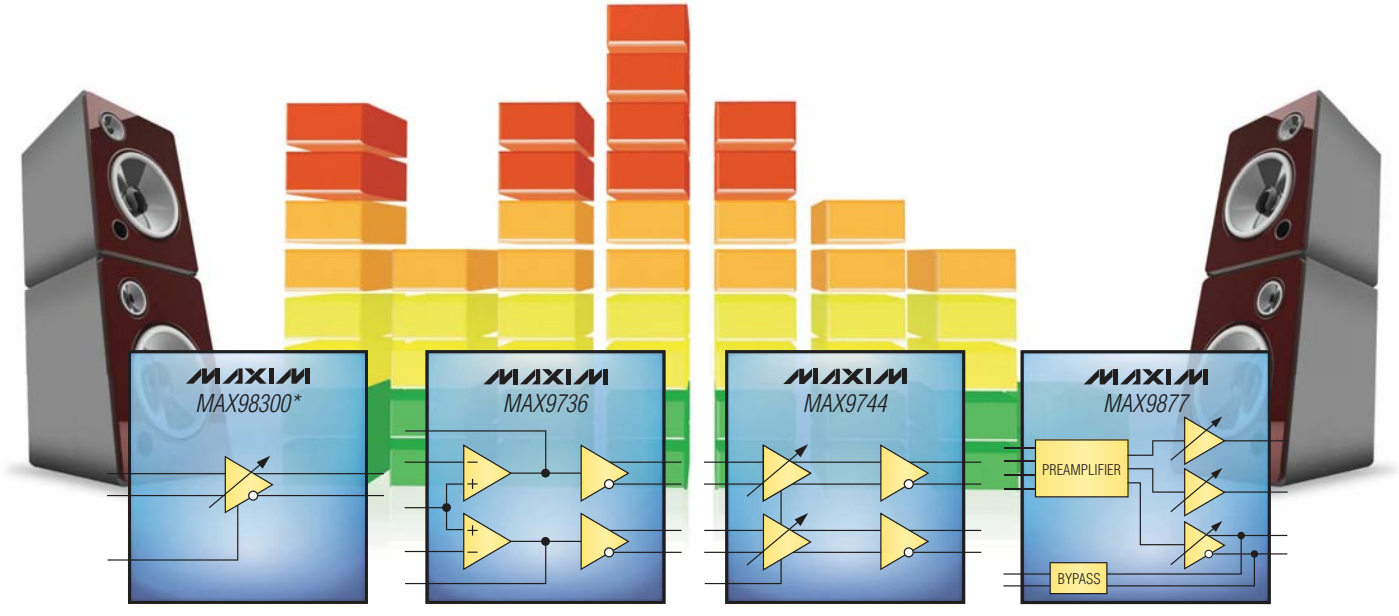

MAX98300*/MAX98302*

- 1 x 2W/2 x 2W
- Filterless up to 60cm with no suppression components
- Industry-low I_Q (1.1mA at 5V)
- UCSP™, TDFN, and TQFN packages
- Applications: netbooks, picture frames, and PNDs

MAX9736/MAX9737

- (2 x 15W, 2 x 6W)/1 x 6W
- Ultra-low EMI—filterless with up to 1.5m cable
- 8V to 28V single supply
- Built-in op amps and simpler BOM reduce cost
- Applications: iPod® docks and TVs

MAX9744/MAX9768

- 2 x 20W/1 x 10W
- Filterless with up to 1m cable
- Clickless analog or I²C volume control
- 4.5V to 14V single supply
- Applications: boomboxes and conference phones

MAX9877/MAX9879

- 2 x 930mW
- Best-in-class RF immunity
- DirectDrive® headphone amp
- Integrated bypass switch
- Space-saving 20-bump UCSP package
- Applications: cell phones

UCSP is a trademark and DirectDrive is a registered trademark of Maxim Integrated Products, Inc.
iPod is a registered trademark of Apple, Inc.
*Future product—contact factory for availability.

www.maxim-ic.com/amplifier-family

MAXIM DIRECT™ www.maxim-ic.com/shop

AVNET electronics marketing www.em.avnet.com/maxim

MAXIM INNOVATION DELIVERED®

For free samples or technical support, visit our website.

eficial effect because it compensates for the decreasing radiance of the LEDs as temperature increases.

Drops in radiance are approximately –0.27%/°C for the blue LED and about –0.35%/°C for the green LED. The radiance of these two LEDs, which are both indium-gallium-nitride types, thus remains almost constant over ambient temperature. The red LED is an aluminium-indium-gallium-phosphorus type, having a radiance drop of approximately –0.77%/°C, and the current source roughly halves this drop.

The R_0 resistors force the logic inputs to logic zero at manual control by connecting or not connecting the IN_1 and IN_2 control leads to V_{DD}, the power-supply voltage. The maximum current flowing through the LEDs, about 26 mA, is far below the nominal current of 350 mA that Avago Technologies (www.avagotech.com) rates for the ASMT-MT00 power RGB (red/green/blue) LED that this circuit uses.

The radiance is sufficient, yet the junction temperature of the LEDs is low. Junction-to-pin thermal resistance for the green LED is 20°C/W. IC_1 dissipates approximately 0.1W. Therefore, you can estimate the junction temperature to be higher than the ambient temperature by less than 2°C (**Reference 1**). Consequently, you increase the LED's expected lifetime well beyond thousands of hours. EDN

REFERENCE

1 Oon, Siang Ling, "The Latest LED Technology Improvement in Thermal Characteristics and Reliability: Avago's Moonstone 3-in-1 RGB High Power LED," White Paper AV02-1752EN, Avago Technologies, Jan 20, 2009, www.avagotech.com/docs/AV02-1752EN.

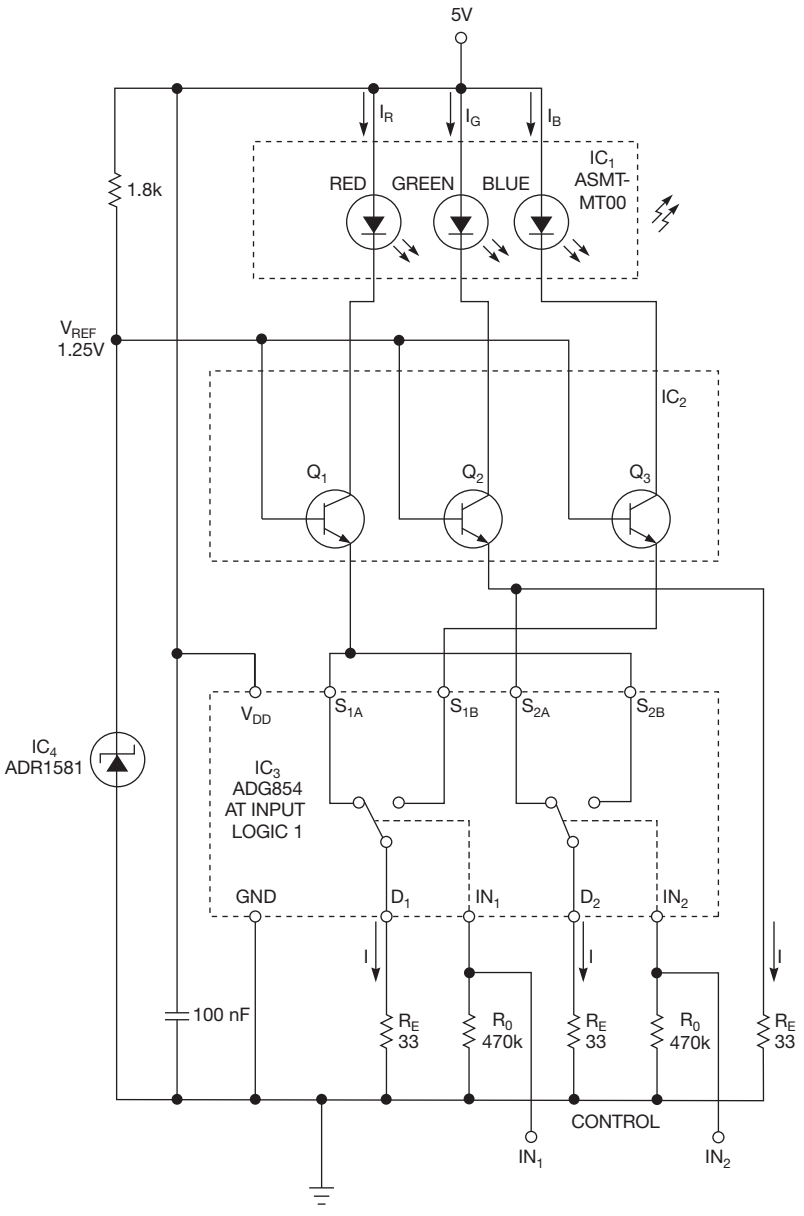

Figure 1 A three-color LED IC emits mixtures of two or three spectrally "pure" colors. The human eye perceives the mixtures as special colors.

Implement trapezoidal velocity profiles in software

Adil Ansari, General Electric, Energy Division, Atlanta, GA

Trapezoidal velocity profiles provide smooth motion for starting and stopping motor-control systems. **Figure 1** shows a velocity-profile section to be implemented in software that you can use to provide digital or analog control to a motor. In the **figure**, $\dot{\theta}$ represents the desired motor velocity (trapezoidal velocity profile) and $\dot{\theta}_{MAX}$ represents the maximum motor velocity.

You predetermine the velocity profile based on the load acceleration and deceleration requirements. A trapezoidal velocity profile has three regions: accel-

eration (Interval 1), constant velocity (Interval 2), and deceleration (Interval 3). **Equation 1** shows the calculation of the desired position, θ_D, or the distance the motor needs to travel, as the area under the curve of **Figure 1**, which is an integral of the velocity.

$$\theta_D = \underbrace{\int_{T1}^{T2} \frac{\dot{\theta}_{MAX} \times t}{(T2-T1)}\,dt}_{\text{INTERVAL 1}} + \underbrace{\int_{T2}^{T3} \dot{\theta}_{MAX}\,dt}_{\text{INTERVAL 2}} + \underbrace{\int_{T3}^{T4} \frac{\dot{\theta}_{MAX} \times t}{(T4-T3)}\,dt}_{\text{INTERVAL 3}}, \quad (1)$$

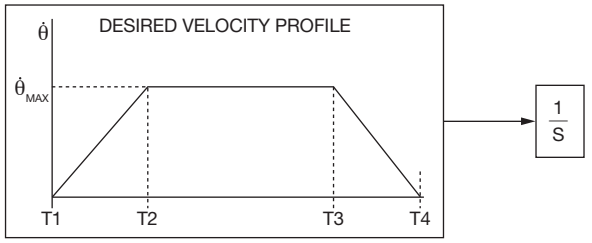

Figure 1 You can implement this velocity-profile section in software that you can use to provide digital or analog motor control.

LISTING 1 PSEUDO CODE

```
// Interval I
// T is sampling interval

    Tint1 = T2 – T1  // Interval I
          For  t = 0 .. Tint1
                       {
                          θi =A.t²
                           t = t + T
                       }

// Interval II

    Tint2 = T3 – T2  // Interval II
          For  t = 0 .. Tint2
                      {
                          θi = θi +B.t
                           t = t + T
                      }

// Interval III

    Tint3 = T4 – T3  // Interval III
          For  t = 0 .. Tint3
          {
            θi = θi+ D – C.(Tint3 – t)²
             t = t + T
          }
```

TRAPEZOIDAL VELOCITY PROFILES PROVIDE SMOOTH MOTION FOR STARTING AND STOPPING MOTOR-CONTROL SYSTEMS.

In **Equation 1**, each integral term represents the areas under the curve for regions 1, 2, and 3, respectively. To implement the velocity profile in a position-control loop, the software calculates the position input at every sampling period, T, by a point-by-point numerical integration of $\dot{\theta}$ (**Listing 1**). You can download the code for a Matlab simulation of this algorithm from the online version of this Design Idea at www.edn.com/100812dib.

Equations 2, **3**, and **4** give the value of $\dot{\theta}$ for intervals 1, 2, and 3, respectively:

$$\text{INTERVAL 1: } \dot{\theta} = \frac{\dot{\theta}_{MAX} \times t}{(T2-T1)} t \quad \text{FOR } t = T1 \text{ TO } T2. \quad (2)$$

$$\text{INTERVAL 2: } \dot{\theta} = \dot{\theta}_{MAX} \quad \text{FOR } t = T2 \text{ TO } T3. \quad (3)$$

$$\text{INTERVAL 3: } \dot{\theta} = \frac{\dot{\theta}_{MAX}}{(T4-T3)} t \quad \text{FOR } t = T3 \text{ TO } T4. \quad (4)$$

Next, calculate the position input to the control loop θ_i every sampling period for each of the intervals by numerically integrating $\dot{\theta}$ in **equations 2**, **3**, and **4**:

$$\text{INTERVAL 1: } \theta_i = \frac{\dot{\theta}_{MAX}}{2(T2-T1)} \times t^2 \,|\, (t = iT,\ i = 0..(T2-T1)). \quad (5)$$

$$\text{INTERVAL 2: } \theta_i = \dot{\theta}_{MAX} \times (t-T2) \,|\, (t = iT,\ i = 0..(T3-T2)). \quad (6)$$

For Interval 3, you calculate the area under the curve (shaded region in **Figure 2**) as t goes from T3 to T4 by subtracting the area of the triangle for the interval [t .. T4] from the area of the entire triangle for the interval [T3 .. T4]:

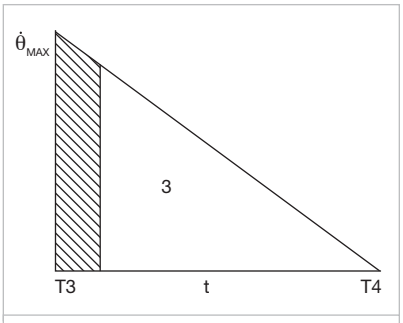

Figure 2 You use subtraction to calculate the area under the curve (shaded region) as t goes from T3 to T4.

$$\text{INTERVAL 3: } \theta_i = \left[\frac{(T4-T3)\,\dot{\theta}_{MAX}}{2} - \frac{\dot{\theta}_{MAX}}{2(T4-T3)} \times (T4-T3-t)^2 \,|\, (t = iT,\ i = 0..t..(T4-T3)) \right]. \quad (7)$$

To simplify the implementation, you can define constants A, B, C, and D (D equals the area under the curve of Interval 3) such that

$$A = \frac{\dot{\theta}_{MAX}}{2(T2-T1)};\ B = \dot{\theta}_{MAX};\ C = \frac{\dot{\theta}_{MAX}}{2(T4-T3)};\ D = \frac{T4-T3}{2} \times \dot{\theta}_{MAX}. \quad (8)$$

$$\theta_i = A \times t^2\big|_{(t = iT,\ i = 0..(T2-T1))} + B \times t\big|_{(t = iT,\ i = 0..(T3-T2))} + (D - C \times (T4 - T3 - t)^2)\big|_{(t = iT,\ i = 0..(T4-T3))}. \quad (9)$$

You can then express θ_i as **Equation 9** (above) shows. EDN

ROHS recast: Electronics industry braces for further regulation

The EU (European Union) is recasting its ROHS (restriction-of-hazardous-substances) environmental directive. By the end of this year or early next year, a ROHS recast will likely carve away exclusions and possibly add new substances to the current list of six banned substances.

The European Parliament has postponed its final vote on the new directive until October because the EU's Environment Committee is meeting with the Council of Ministers to collaborate on the expansion of ROHS. The ROHS recast will include medical and monitoring products, which have begun conversion to ROHS compliance.

At Arrow Electronics Inc (www.arrow.com), the components division has prepared for additional ROHS substances and fewer exclusions. "With more end equipment coming into scope, we find more of our customers needing to comply for the first time, and, therefore, they're looking to Arrow for help in … [achieving] … compliance," says Peter Kong, president of Arrow Electronics Global Components.

The Environment Committee took a pass on adding PVC (polyvinyl chloride) and BFRs (brominated flame retardants) to the original six banned substances. Instead, it asked for further study. The electronics industry, however, is expecting the committee to eventually add these substances. Given that expectation, an industry coalition voluntarily stopped using PVC and BFRs, and it has asked the industry to follow suit in finding alternatives. The group includes Acer (www.acer.com), Dell (www.dell.com), Hewlett-Packard (www.hp.com), and Sony Ericsson (www.sonyericsson.com), as well as the environmental groups the International ChemSec (Chemical Secretariat, www.chemsec.org), Clean Production Action (www.cleanproduction.org), and the European Environmental Bureau (www.eeb.org).

Many electronics companies have voluntarily moved to eliminate PVC and BFRs from their products. Some industry watchers are skeptical, however. "These companies are acting the way they do because Greenpeace [www.greenpeace.org] has a gun to their heads," says Fern Abrams, director of government relations and environmental policy at the IPC (Institute for Interconnecting and Packaging Electronic Circuits) Association Connecting Electronics Industries (www.ipc.org). "I would suspect this was a business decision; it's easier to negotiate with than against ChemSec." Abrams doubts that the companies will eliminate all BFRs and PVC substances. "The footnotes to their statements on their Web sites … look like credit-card agreements," she says.

"We've worked on the educational aspect of this issue," says Alexandra McPherson, managing partner of Clean Production Action. "We've produced a technical report that shows the feasibility of the transition away from these substances. We did the report both for the policy members in Europe who will decide [what ROHS will include] and for the electronics supply chain so we could tell manufacturers what strategies companies have developed to overcome the challenges and barriers to removing PVC and BFR" (**Reference 1**).

One of the impending restrictions that strikes terror in the heart of the electronics supply chain is the almost-inevitable prospect that ROHS will become a CE-mark directive—a significant event, according to Gary Nevison (**photo**), legislation and environmental affairs manager at the UK-based distributor Premier Farnell (www.premierfarnell.com) and its sister company Newark (www.newark.com) in the United States. "Companies will have to produce significant documentation. Small and midsized companies are terrified of this [step]," he says, which would require a significant escalation in reporting of content in electronic products.

"It will be an interesting awakening, too, for some companies," says Ken Stanvick, senior vice president and co-founder of Design Chain Associates (www.designchainassociates.com). "The idea of having technical documentation to defend your claim as ROHS-compliant will be difficult."

In response to growing regulations, IPC has launched a campaign asking regulating bodies such as the European Parliament to make sure that they base their regulations on scientific facts. "When you look at some of the proposed amendments, you begin to abandon hope that it will be a scientific process," says the IPC's Abrams. "Some of the amendments are disturbing to anyone who cares about science."—**by Rob Spiegel**

REFERENCE

1 Nimpuno, Nardano; Alexandra McPherson; and Tanvir Sadique, "Greening Consumer Electronics—moving away from bromine and chlorine," ChemSec and Clean Production Action, September 2009, http://cleanproduction.org/library/GreeningConsumerElectronics.pdf.

The New Resource Dedicated Exclusively to Designers and Developers of Medical Electronic Devices

DesignMED

Emerging Technologies for Medical Electronic Devices

CONFERENCE & EXPOSITION

February 8-10, 2011

Anaheim Convention Center
Anaheim, CA

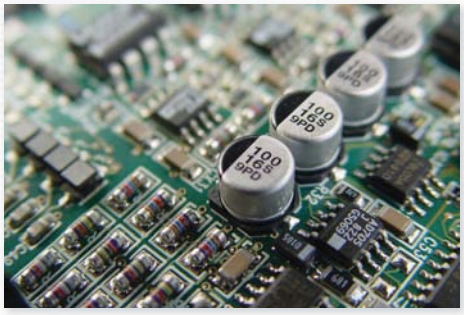

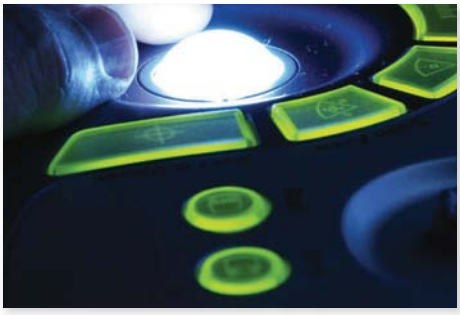

Meet hundreds of the nation's leading suppliers presenting the latest in emerging technologies that can be incorporated into the next generation of medical electronics, including:

- Analog ICs
- Applied Systems
- Communications & Networking Devices
- Components
- DSPs
- EDAs
- Graphics/Graphics-Interface ICs
- Memory & Storage
- Microcontrollers & Processors
- Power
- Processors & Tools
- Programmable ICs
- RF
- Semiconductor Processes & IP
- Software
- Test & Measurement...*and much more!*

Supported by:

MED Medical Electronics Design · EDN Electronics Design, Strategy, News · design news

Co located with:

Medical Design & Manufacturing M&M WEST · Electronics WEST

Produced and managed by:
Canon Communications LLC
cancom.com

Your Project Acceleration RESOURCE

For more information on attending or exhibiting, please visit

DesignMedShow.com

Submit your registration request online today. Please use Promo Code: AA

16212_W_DMD11

productroundup

DISCRETE SEMICONDUCTORS

RF transistors feature ESD protection

The BFP640ESD, BFP720ESD, and BFP740ESD line of RF transistors includes ESD protection as great as 2-kV human-body model for reliable, high-sensitivity wireless-communication devices. The transistors can act as the low-noise-amplifier stage of the RF-signal chain, reducing the risk of ESD that can lead to failure or damage to the wireless system. Maximum input power is 20 dBm, and the noise figure is 0.6 dB at 2.4 GHz. Prices for sampling units range from 40 to 50 cents each.

Infineon Technologies, www.infineon.com

LED-driver IC increases reliability of HB lamps

The multitopology ZXLD1370 LED-driver IC increases the reliability of HB lamps in automotive, industrial, and commercial lighting systems. It operates in buck, boost, and buck-boost modes and combines patent-pending control loops and high-side current sensing to ensure accurate current control of LED strings. With a 6 to 60V operating input-voltage range and a typical 1% output-current tolerance in all topologies, the driver delivers the high current levels and tight interlamp luminance matching that HB LEDs require. The driver achieves active thermal management through a dedicated external input and fault-diagnosis outputs that report the status of the LED driver and load. The driver works with an external MOSFET and has 1-MHz switching frequency with a 100-kHz low-frequency clamp, reducing the size of the external inductor. The ZXLD1370 sells for $1 (1000).

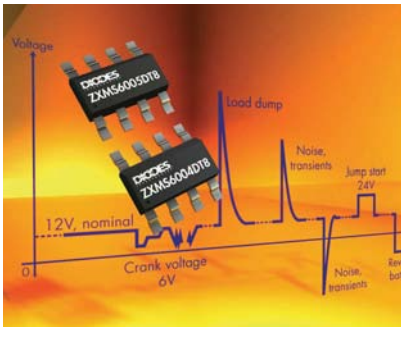

Diodes Inc, www.diodes.com

Small-footprint MOSFET targets portable-system applications

This MicroFET MOSFET series comes in a 1.6×1.6×0.55-mm footprint for battery-charging, dc/dc-conversion, level-shifted-load-switching, and boost-switching applications.

The line includes common topology choices, such as single-P-channel/Schottky-diode combo, single-N-channel/Schottky-diode combo, dual and single P-channel, dual and single N-channel, and complementary pair. Prices start at 33 cents (1000).

Fairchild Semiconductor, www.fairchildsemi.com

Power-MOSFET family operates from –30 to +100V

The IRMLx family of SOT-23-packaged power MOSFETs have voltages ranging from –30 to +100V. Maximum drain current ranges from

3.6 to 1.6A, and gate charge ranges from 4.8 to 2.5 nC. Typical on-resistances range from 51 to 178 mΩ at 10V and 82 to 190 mΩ at 4.5V. Maximum on-resistances range from 64 to 220 mΩ at 10V and 103 to 235 mΩ at 4.5V. Prices begin at 10 cents (10,000).

International Rectifier, www.irf.com

30V N-channel power MOSFETs integrate Schottky diodes

The NTMFS4897NF, NTMFS-4898NF, and NTMFS4899NF series of N-channel power MOSFETs have maximum on-resistances of 2, 3, and 5 mΩ, respectively, at 10V, targeting use as the synchronous side in buck-converter applications to achieve high power efficiency. Typical gate-charge specs are 39.6, 25.6, and 12.2 nC, respectively, at a gate-to-source voltage of 4.5V. The devices come in 536-mm SO-8FL packages; prices range from 45 to 90 cents (10,000).

On Semiconductor, www.onsemi.com

N-channel power MOSFETs feature low on-resistance

The SiHP12N50C-E3, SiHF-12N50C-E3, and SiHF12N50C-E3 family of N-channel power MOSFETs comes in TO-220, TO-220 Fullpack, and D²Pak packages, respectively. The 500V, 12A devices feature 0.555Ω on-resistances at a 10V gate drive and a gate charge of 48 nC. Gate charge times on-resistance, a key figure of merit for MOSFETs targeting power-

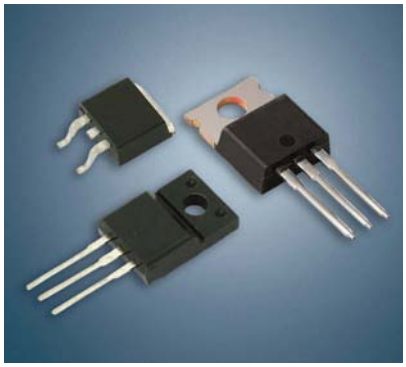

conversion applications, is 26.64Ω-nC. Prices for the SiHP12N50C-E3, SiHF-12N50C-E3, and SiHF12N50C-E3 start at 65, 71, and 72 cents (sample quantities), respectively.

Vishay Siliconix, www.vishay.com

SWITCHES AND RELAYS

AC switches provide enhanced protection

This line of ac switches targets use in appliance-motor control and eliminates the need for additional components to ensure smooth switching and provide protection. The ACST4, ACST6, and ACST8 switches include a power TRIAC delivering enhanced commutation performance to turn off the motor without requiring external components to suppress voltage transients. Unlike conventional TRIACs, the switches integrate protection against surge voltages as high as 2 kV on the ac line. The ACST410 and ACST610 switches have a trigger current of 10 mA, allowing direct control by a CMOS device. The ACST435 and ACST830 switches specify noise immunity of at least 1000 and 2000V/µsec, respectively. The families comprise 14 part numbers, and prices start at 70 cents (10,000).

STMicroelectronics, www.st.com

EDN ADVERTISER INDEX

EDN provides this index as an additional service. The publisher assumes no liability for errors or omissions.

Pen puzzle pauses printer

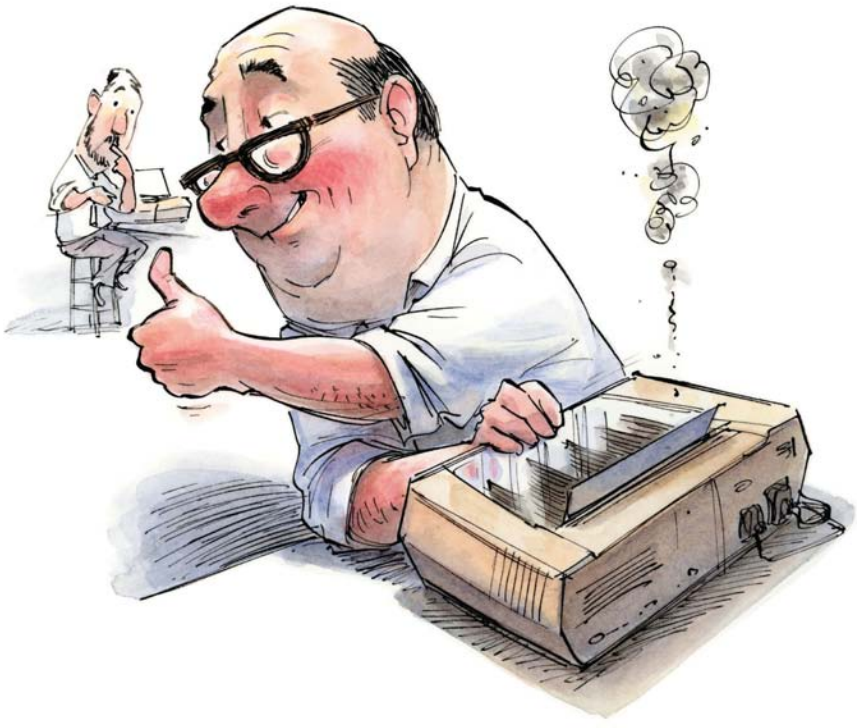

DANIEL VASCONCELLOS

Many years ago, I was working for a company that designed and manufactured desktop ink-jet printers. We had just released a model that used a new higher-resolution color pen, but we were having high infant mortality on the new printers. Management asked me to investigate the failures that were primarily occurring in the head drive's ASIC, which fired the pen. The ASIC contained switching FETs controlling the firing currents to the pens' nozzle resistors, which expelled the ink from the nozzles. The ASIC also included circuitry that could identify the type of pen installed. We uncapped some of the failed ASICs to expose the raw die and found that the damage always occurred on one nozzle-driver FET, which had suffered a catastrophic failure.

Coincidentally, I found that the pen engineers were puzzling over a rash of warranty returns that involved burnt traces on the same line that the failed FET would drive. In the earlier-model pens, these traces went to ground. In the new generation, the firing signals went to a nozzle's firing resistor. Even though the FETs were current-limited, they wouldn't withstand driving a hard short for long.

It looked as though we had installed the "wrong" pens in the printers and that the head drive's ASIC had tried to fire a signal line that was tied to ground, resulting in the failure of both the head-drive ASIC and the pen's ground traces. However, that scenario seemed impossible because, when firmware recognized the installation of a new pen, it first verified that it was the proper type. If it had been the wrong pen, it would have flashed an error to the user and refused to energize the pen. The system was somehow energizing the pen anyway; nothing else made sense. How could this problem have occurred?

First, I verified that the firmware would indeed reject the incorrect pen. Normal pen-changing procedure involved opening an access door, removing the old pen, installing the new one, and then closing the door. A door sensor let the printer know that the door was open and that a pen change was possible. The firmware then read ID bits on the pen and, once it verified them as correct, ran a warm-up routine that involved firing all the nozzles. Initially, I could not get the printer to accept the wrong pen, although I tried variations of cycling power and inserting and removing various pens in an attempt to confuse it. I then remembered my mother's warning that, to avoid risk of shock, I should unplug the toaster before attempting to remove stuck toast. What if a printer user thought in the same way: Unplug an electrical device before reaching inside it?

To test this idea, I opened the access door and unplugged the printer. I then swapped out the old pen for the wrong pen, closed the access door, and plugged in the printer. It worked perfectly. The printer accepted the new pen and immediately began trying to warm it up to operating temperatures by firing all the nozzles. With the wrong pen, this attempt meant driving one of the nozzle-firing FETs into a hard short circuit. It quickly blew up both the head drive's ASIC and the pen.

Failure analysis revealed that the failures were identical to those that had occurred in the field. But why did the printer accept the wrong pen? We investigated the firmware and found that the printer performed its pen-ID check only after someone opened the door. No one had figured that a user could get a new pen into the printer without the printer's "knowledge." We had to live with the failures until a new version of firmware could make it out to the field, at which time the failures went away. EDN

Steve Soar is retired from a career as an analog, motor-drive, ASIC-development, servo, optical-sensor, motion-sensor, and bottle-washer engineer at Hewlett-Packard.

When high performance and low power make the design, ADI converters make the difference.

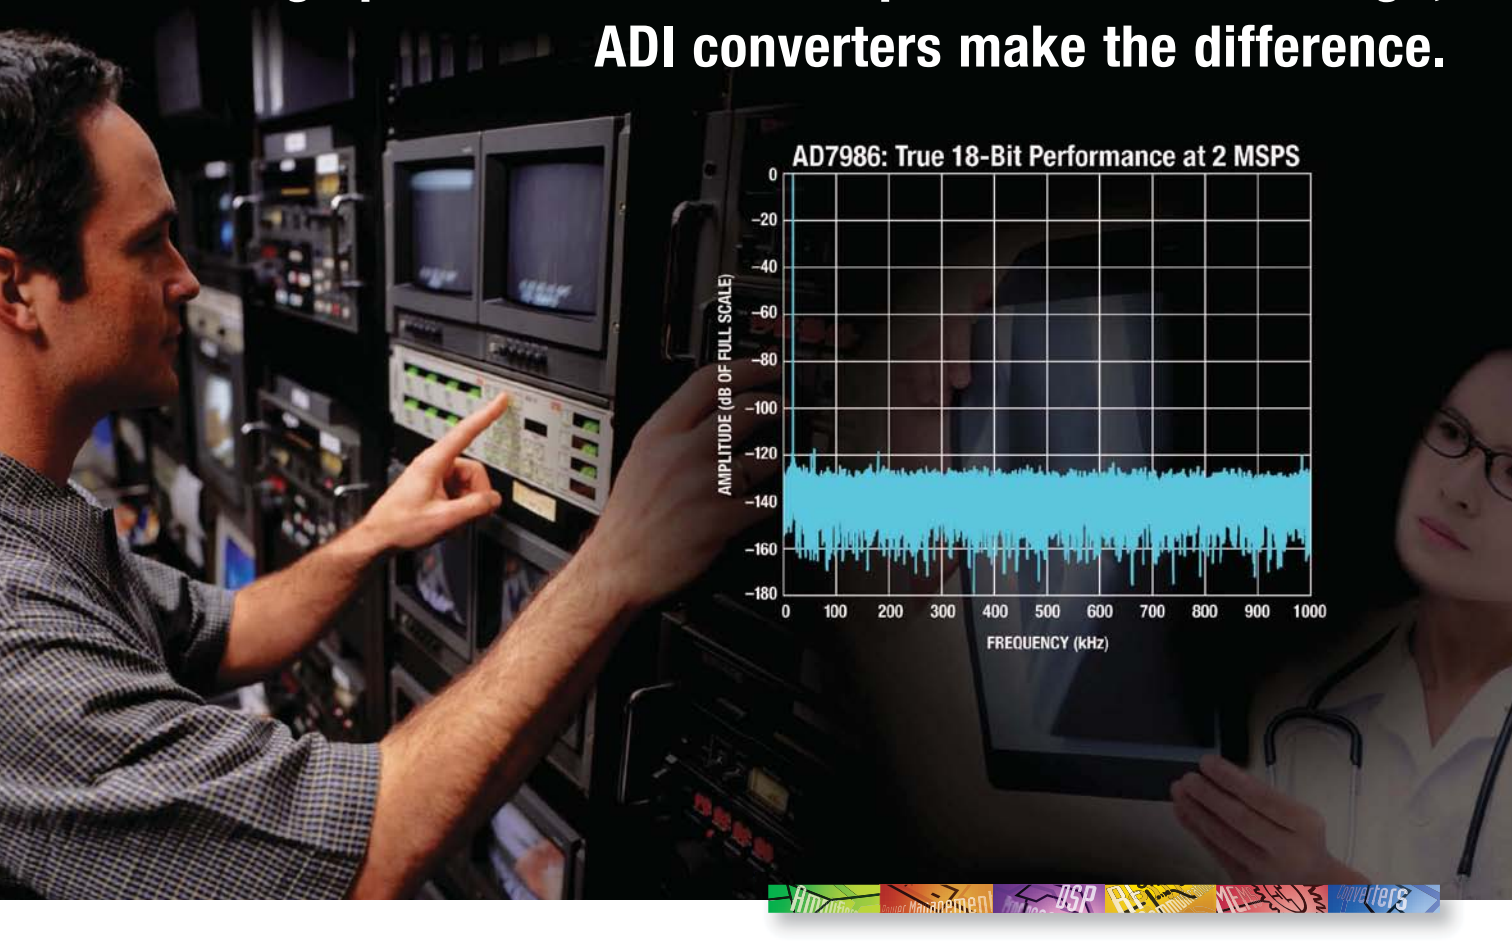

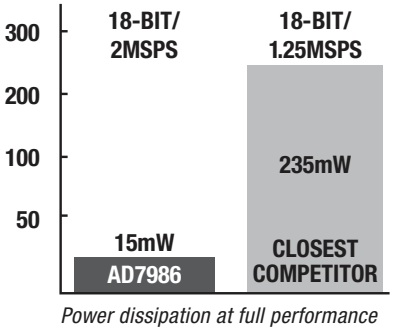

Power dissipation at full performance

15 mW, 18-bit, 2 MSPS, ADC delivers industry-leading performance without the heat.

Designers of high channel count, high precision data acquisition systems inevitably face the challenges of minimizing both power consumption and design footprint. Achieving maximum system throughput under these circumstances required either multiplexed high speed resolution ADCs, which generate more heat, or multiple low speed resolution ADCs, which take up extra board space. With power dissipation levels nearly 15× lower than the competition and a smaller footprint, the AD7986 eliminates the need for such trade-offs. And you'll find many more industry-leading examples like this one throughout our high precision SAR converter portfolio at *www.analog.com/PulSAR.*

Part Number	Speed (MSPS)	Resolution (Bits)	Features	Power Dissipation (mW)
AD7986	2	18	97 dB SNR, ±2.5 LSB INL	15
AD7985	2.5	16	90 dB SNR, ±1.5 LSB INL	15.5
AD7944	2.5	14	84.5 dB SNR, ±1 LSB INL	15.5

All products available in 20-lead LFCSP package.

MAKEADIFFERENCE

www.analog.com

ANALOG DEVICES

Data acquisition just got a lot easier.

More ways to control. More ways to connect.

The new Agilent 34972A Data Acquisition Switch Unit takes our best-selling Agilent 34970A to the next level. For starters, you get convenient built-in LAN and USB connectivity. Plus, you can control your data acquisition remotely via Web interface. And transfer logged data to your PC with a simple flash drive. No more expensive adapters and connectors. That's easy. That's Agilent.

NEW 34972A	34970A
USB and LAN	GPIB and RS232
Graphical web interface Benchlink data logger software SCPI programming	Benchlink data logger software SCPI programming
3-slot LXI unit with built-in 6 ½ digit DMM	
~~$1,845~~ $1,597*	$1,597*

Hurry! For a limited time, get the 34972A at the 34970A price.
www.agilent.com/find/Agilent34972A

*Prices are in USD and are subject to change.
See participating distributors for details.

Agilent Technologies

Agilent and our Distributor Network
Right Instrument. Right Expertise. Delivered Right Now.

RS ELECTRONICS
888-852-7234
www.rselectronics.com